



---

Junos<sup>®</sup> OS

## Layer 3 VPNs Feature Guide for Routing Devices



---

Modified: 2018-06-08

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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 Layer 3 VPNs Feature Guide for Routing Devices*  
Copyright © 2018 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 <https://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 . . . . .	xxvii
	Documentation and Release Notes . . . . .	xxvii
	Supported Platforms . . . . .	xxvii
	Using the Examples in This Manual . . . . .	xxvii
	Merging a Full Example . . . . .	xxviii
	Merging a Snippet . . . . .	xxviii
	Documentation Conventions . . . . .	xxix
	Documentation Feedback . . . . .	xxxi
	Requesting Technical Support . . . . .	xxxi
	Self-Help Online Tools and Resources . . . . .	xxxi
	Opening a Case with JTAC . . . . .	xxxii
<b>Part 1</b>	<b>Common Configuration for All VPNs</b>	
<b>Chapter 1</b>	<b>VPNs Overview . . . . .</b>	<b>3</b>
	VPLS . . . . .	3
	Types of VPNs . . . . .	4
	Layer 2 VPNs . . . . .	4
	Layer 3 VPNs . . . . .	5
	VPLS . . . . .	5
	Virtual-Router Routing Instances . . . . .	6
	VPNs and Logical Systems . . . . .	7
	Layer 2 VPNs . . . . .	7
	Routers in a VPN . . . . .	8
<b>Chapter 2</b>	<b>Assigning Routing Instances to VPNs . . . . .</b>	<b>9</b>
	Configuring Routing Instances on PE Routers in VPNs . . . . .	9
	Configuring the Routing Instance Name for a VPN . . . . .	10
	Configuring the Description . . . . .	10
	Configuring the Instance Type . . . . .	10
	Configuring Interfaces for VPN Routing . . . . .	11
	General Configuration for VPN Routing . . . . .	11
	Configuring Interfaces for Layer 3 VPNs . . . . .	12
	Configuring Interfaces for Carrier-of-Carriers VPNs . . . . .	12
	Configuring Unicast RPF on VPN Interfaces . . . . .	13
	Configuring the Route Distinguisher . . . . .	13
	Configuring Automatic Route Distinguishers . . . . .	14
	Configuring Virtual-Router Routing Instances in VPNs . . . . .	14
	Configuring a Routing Protocol Between the Service Provider Routers . . . . .	15
	Configuring Logical Interfaces Between Participating Routers . . . . .	15

	Configuring Path MTU Checks for VPN Routing Instances . . . . .	16
	Enabling Path MTU Checks for a VPN Routing Instance . . . . .	17
	Assigning an IP Address to the VPN Routing Instance . . . . .	17
<b>Chapter 3</b>	<b>Distributing Routes in VPNs . . . . .</b>	<b>19</b>
	Enabling Routing Information Exchange for VPNs . . . . .	19
	Configuring IBGP Sessions Between PE Routers in VPNs . . . . .	19
	Configuring Aggregate Labels for VPNs . . . . .	21
	Configuring a Signaling Protocol and LSPs for VPNs . . . . .	22
	Using LDP for VPN Signaling . . . . .	22
	Using RSVP for VPN Signaling . . . . .	24
	Configuring Policies for the VRF Table on PE Routers in VPNs . . . . .	25
	Configuring the Route Target . . . . .	26
	Configuring the Route Origin . . . . .	26
	Configuring an Import Policy for the PE Router's VRF Table . . . . .	27
	Configuring an Export Policy for the PE Router's VRF Table . . . . .	29
	Applying Both the VRF Export and the BGP Export Policies . . . . .	30
	Configuring a VRF Target . . . . .	31
	Configuring the Route Origin for VPNs . . . . .	32
	Configuring the Site of Origin Community on CE Router A . . . . .	33
	Configuring the Community on CE Router A . . . . .	33
	Applying the Policy Statement on CE Router A . . . . .	33
	Configuring the Policy on PE Router D . . . . .	34
	Configuring the Community on PE Router D . . . . .	34
	Applying the Policy on PE Router D . . . . .	35
<b>Chapter 4</b>	<b>Distributing VPN Routes with Target Filtering . . . . .</b>	<b>37</b>
	Configuring BGP Route Target Filtering for VPNs . . . . .	37
	BGP Route Target Filtering Overview . . . . .	37
	Configuring BGP Route Target Filtering for VPNs . . . . .	38
	Example: BGP Route Target Filtering for VPNs . . . . .	39
	Example: Configuring BGP Route Target Filtering for VPNs . . . . .	41
	Configure BGP Route Target Filtering on Router PE1 . . . . .	42
	Configure BGP Route Target Filtering on Router PE2 . . . . .	43
	Configure BGP Route Target Filtering on the Route Reflector . . . . .	46
	Configure BGP Route Target Filtering on Router PE3 . . . . .	47
	Configuring Static Route Target Filtering for VPNs . . . . .	50
	Understanding Proxy BGP Route Target Filtering for VPNs . . . . .	50
	Example: Configuring Proxy BGP Route Target Filtering for VPNs . . . . .	51
	Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs . . . . .	67
	Reducing Network Resource Use with Static Route Target Filtering for VPNs . . .	84
<b>Chapter 5</b>	<b>Configuring Forwarding Options for VPNs . . . . .</b>	<b>85</b>
	Chained Composite Next Hops for Layer 2 VPNs and Layer 2 Circuits . . . . .	85
	Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs . . . . .	86



<b>Chapter 6</b>	<b>Configuring Graceful Restart for VPNs</b>	<b>101</b>
	VPN Graceful Restart	101
	Configuring Graceful Restart for VPNs	102
	Enabling Unicast Reverse-Path Forwarding Check for VPNs	104
	Example: Configuring Unicast Reverse-Path-Forwarding Check	105
	Understanding How Unicast Reverse Path Forwarding Prevents Spoofed	
	IP Packet Forwarding	105
	Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent	
	DoS and DDoS Attacks	105
<b>Chapter 7</b>	<b>Configuring Class of Service for VPNs</b>	<b>115</b>
	VPNs and Class of Service	115
	Rewriting Class of Service Markers and VPNs	115
<b>Chapter 8</b>	<b>Pinging VPNs</b>	<b>117</b>
	Pinging VPNs, VPLS, and Layer 2 Circuits	117
	Setting the Forwarding Class of the Ping Packets	118
	Pinging a VPLS Routing Instance	118
	Pinging a Layer 2 VPN	119
	Pinging a Layer 3 VPN	119
	Pinging a Layer 2 Circuit	119
<b>Part 2</b>	<b>Common Configuration for Layer 3 VPNs</b>	
<b>Chapter 9</b>	<b>Layer 3 VPN Overview</b>	<b>123</b>
	Understanding Layer 3 VPNs	123
	Components of a Layer 3 VPN	124
	Layer 3 VPN Terminology	124
	Layer 3 VPN Architecture	125
	Understanding Layer 3 VPN Forwarding Through the Core	127
	Understanding Layer 3 VPN Attributes	129
	Supported Layer 3 VPN Standards	130
	Introduction to Configuring Layer 3 VPNs	130
	Layer 3 VPNs for IPv4 and IPv6 Overview	133
<b>Chapter 10</b>	<b>Assigning Routing Instances to Layer 3 VPNs</b>	<b>137</b>
	Understanding Routing Instances for Layer 3 VPNs	137
	Configuring Logical Units on the Loopback Interface for Routing Instances in	
	Layer 3 VPNs	137
<b>Chapter 11</b>	<b>Creating Unique VPN Routes Using VRF Tables</b>	<b>141</b>
	Understanding Virtual Routing and Forwarding Tables	141
	Understanding VRF Localization in Layer 3 VPNs	145
	Maximizing VPN Routes Using VRF Localization for Layer 3 VPNs	146
	Example: Improving Scalability Using VRF Localization for Layer 3 VPNs	147
	Filtering Packets in Layer 3 VPNs Based on IP Headers	159
	Egress Filtering Options	160
	Support on Aggregated and VLAN Interfaces for IP-Based Filtering	161
	Support on ATM and Frame Relay Interfaces for IP-Based Filtering	161

	Support on Ethernet, SONET/SDH, and T1/T3/E3 Interfaces for IP-Based Filtering . . . . .	162
	Support on SONET/SDH and DS3/E3 Channelized Enhanced Intelligent Queuing Interfaces for IP-Based Filtering . . . . .	162
	Support on Multilink PPP and Multilink Frame Relay Interfaces for IP-Based Filtering . . . . .	164
	Support for IP-Based Filtering of Packets with Null Top Labels . . . . .	164
	General Limitations on IP-Based Filtering . . . . .	165
	Configuring a Label Allocation and Substitution Policy for VPNs . . . . .	166
<b>Chapter 12</b>	<b>Distributing Routes in Layer 3 VPNs . . . . .</b>	<b>169</b>
	Configuring Routing Between PE and CE Routers in Layer 3 VPNs . . . . .	169
	Configuring BGP Between the PE and CE Routers . . . . .	169
	Configuring OSPF Between the PE and CE Routers . . . . .	170
	Configuring OSPF Version 2 Between the PE and CE Routers . . . . .	170
	Configuring OSPF Version 3 Between the PE and CE Routers . . . . .	171
	Configuring OSPF Sham Links for Layer 3 VPNs . . . . .	171
	OSPF Sham Links Overview . . . . .	171
	Configuring OSPF Sham Links . . . . .	172
	OSPF Sham Links Example . . . . .	173
	Configuring an OSPF Domain ID . . . . .	174
	Hub-and-Spoke Layer 3 VPNs and OSPF Domain IDs . . . . .	176
	Configuring RIP Between the PE and CE Routers . . . . .	177
	Configuring Static Routes Between the PE and CE Routers . . . . .	179
	Configuring an OSPF Domain ID for a Layer 3 VPN . . . . .	179
	Configuring Interfaces on Router PE1 . . . . .	180
	Configuring Routing Options on Router PE1 . . . . .	181
	Configuring Protocols on Router PE1 . . . . .	181
	Configuring Policy Options on Router PE1 . . . . .	181
	Configuring the Routing Instance on Router PE1 . . . . .	182
	Configuration Summary for Router PE1 . . . . .	183
	Example: Giving VPN Backbones Priority with OSPFv2 Sham Links for Layer 3 VPNs . . . . .	185
	OSPFv2 Sham Links Overview . . . . .	185
	Example: Configuring OSPFv2 Sham Links . . . . .	186
	Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs . . . . .	195
	Configuring an LDP-over-RSVP VPN Topology . . . . .	196
	Enabling an IGP on the PE and P Routers . . . . .	199
	Enabling LDP on the PE and P Routers . . . . .	199
	Enabling RSVP and MPLS on the P Router . . . . .	201
	Configuring the MPLS LSP Tunnel Between the P Routers . . . . .	201
	Configuring IBGP on the PE Routers . . . . .	202
	Configuring Routing Instances for VPNs on the PE Routers . . . . .	203
	Configuring VPN Policy on the PE Routers . . . . .	204
	LDP-over-RSVP VPN Configuration Summarized by Router . . . . .	206
	Router PE1 . . . . .	206
	Router P1 . . . . .	207
	Router P2 . . . . .	208

	Router P3 . . . . .	208
	Router PE2 . . . . .	209
<b>Chapter 13</b>	<b>Associating Interfaces with Layer 3 VPNs . . . . .</b>	<b>211</b>
	Configuring an Application-Based Layer 3 VPN Topology . . . . .	211
	Configuration on Router A . . . . .	212
	Configuration on Router E . . . . .	214
	Configuration on Router F . . . . .	214
<b>Chapter 14</b>	<b>Carrying IPv4 Traffic Over Layer 3 VPNs . . . . .</b>	<b>217</b>
	Understanding VPN-IPv4 Addresses and Route Distinguishers . . . . .	217
	Understanding IPv4 Route Distribution in a Layer 3 VPN . . . . .	221
	Distribution of Routes from CE to PE Routers . . . . .	221
	Distribution of Routes Between PE Routers . . . . .	222
	Distribution of Routes from PE to CE Routers . . . . .	224
	Configuring IPv4 Packet Forwarding for Layer 3 VPNs . . . . .	225
<b>Chapter 15</b>	<b>Carrying IPv6 Traffic Over Layer 3 VPNs . . . . .</b>	<b>227</b>
	Understanding IPv6 Layer 3 VPNs . . . . .	227
	Configuring Layer 3 VPNs to Carry IPv6 Traffic . . . . .	227
	Configuring IPv6 on the PE Router . . . . .	228
	Configuring the Connection Between the PE and CE Routers . . . . .	228
	Configuring BGP on the PE Router to Handle IPv6 Routes . . . . .	228
	Configuring BGP on the PE Router for IPv4 and IPv6 Routes . . . . .	229
	Configuring OSPF Version 3 on the PE Router . . . . .	229
	Configuring Static Routes on the PE Router . . . . .	230
	Configuring IPv6 on the Interfaces . . . . .	230
	Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains . . . . .	231
<b>Chapter 16</b>	<b>Configuring an AS for Layer 3 VPNs . . . . .</b>	<b>241</b>
	Configuring Layer 3 VPNs to Carry IBGP Traffic . . . . .	241
	Example: Configuring a Layer 3 VPN with Route Reflection and AS Override . . . . .	243
	Configuring the Algorithm That Determines the Active Route to Evaluate AS Numbers in AS Paths for VPN Routes . . . . .	253
<b>Chapter 17</b>	<b>Limiting Routes Using Route Resolution and Route Filters . . . . .</b>	<b>255</b>
	Example: Configuring Route Resolution on PE Routers . . . . .	255
	Example: Configuring Route Resolution on Route Reflectors . . . . .	257
	Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs . . . . .	260
<b>Chapter 18</b>	<b>Configuring Protocol Family Qualifiers for Layer 3 VPNs . . . . .</b>	<b>263</b>
	Example: Configuring Layer 3 VPN Protocol Family Qualifiers for Route Filters . . . . .	263
<b>Chapter 19</b>	<b>Configuring Class of Service for Layer 3 VPNs . . . . .</b>	<b>267</b>
	Configuring Traffic Policing in Layer 3 VPNs . . . . .	267
	Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs . . . . .	268

<b>Chapter 20</b>	<b>Enabling Internet Access for Layer 3 VPNs</b>	<b>271</b>
	Non-VRF Internet Access Through Layer 3 VPNs	271
	CE Router Accesses Internet Independently of the PE Router	272
	PE Router Provides Layer 2 Internet Service	272
	Distributed Internet Access Through Layer 3 VPNs	272
	Routing VPN and Internet Traffic Through Different Interfaces for Layer 3	
	VPNs	273
	Configuring Interfaces on Router PE1	274
	Configuring Routing Options on Router PE1	274
	Configuring BGP, IS-IS, and LDP Protocols on Router PE1	274
	Configuring a Routing Instance on Router PE1	275
	Configuring Policy Options on Router PE1	276
	Traffic Routed by Different Interfaces: Configuration Summarized by	
	Router	277
	Router PE1	277
	Routing VPN and Outgoing Internet Traffic Through the Same Interface and	
	Routing Return Internet Traffic Through a Different Interface	279
	Configuration for Router PE1	279
	Routing VPN and Internet Traffic Through the Same Interface Bidirectionally	
	(VPN Has Public Addresses)	280
	Configuring Routing Options on Router PE1	281
	Configuring Routing Protocols on Router PE1	281
	Configuring the Routing Instance on Router PE1	282
	Traffic Routed Through the Same Interface Bidirectionally: Configuration	
	Summarized by Router	283
	Router PE1	283
	Routing VPN and Internet Traffic Through the Same Interface Bidirectionally	
	(VPN Has Private Addresses)	284
	Configuring Routing Options for Router PE1	285
	Configuring a Routing Instance for Router PE1	285
	Configuring Policy Options for Router PE1	286
	Traffic Routed by the Same Interface Bidirectionally (VPN Has Private	
	Addresses): Configuration Summarized by Router	286
	Router PE1	286
	Routing Internet Traffic Through a Separate NAT Device	288
	Centralized Internet Access Through Layer 3 VPNs	296
	Routing Internet Traffic Through a Hub CE Router	296
	Configuring a Routing Instance on Router PE1	297
	Configuring Policy Options on Router PE1	298
	Internet Traffic Routed by a Hub CE Router: Configuration Summarized	
	by Router	299
	Routing Internet Traffic Through Multiple CE Routers	300
	Configuring a Routing Instance on Router PE1	301
	Configuring Policy Options on Router PE1	302
	Configuring a Routing Instance on Router PE3	303
	Configuring Policy Options on Router PE3	303
	Routing Internet Traffic Through Multiple CE Routers: Configuration	
	Summarized by Router	304

<b>Part 3</b>	<b>Configuring Interprovider and Carrier-of-Carrier VPNs</b>	
<b>Chapter 21</b>	<b>Interprovider and Carrier-of-Carrier VPNs Overview</b>	<b>311</b>
	Traditional VPNs, Interprovider VPNs, and Carrier-of-Carriers VPNs	311
	Understanding Interprovider and Carrier-of-Carriers VPNs	312
	Interprovider and Carrier-of-Carrier VPNs Example Terminology	313
	Supported Carrier-of-Carriers and Interprovider VPN Standards	313
	Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics	314
<b>Chapter 22</b>	<b>Configuring Interprovider VPNs</b>	<b>317</b>
	Interprovider VPNs	317
	Linking VRF Tables Between Autonomous Systems	318
	Configuring Next Generation Layer 3 VPNs Options A, B, and C	318
	Configuring Multihop MP-EBGP Between AS Border Routers	319
	Interprovider VPN Example—MP-EBGP Between ISP Peer Routers	319
	Configuration for Router A	320
	Configuration for Router B	321
	Configuration for Router C	322
	Configuration for Router D	323
	Configuration for Router E	324
	Configuration for Router F	326
	Interprovider VPN Example—Multihop MP-EBGP with P Routers	326
	Configuration for Router A	327
	Configuration for Router B	328
	Configuration for Router C	329
	Configuration for Router D	331
	Configuration for Router E	332
	Configuration for Router F	333
	Example: Configuring Interprovider Layer 3 VPN Option A	334
	Example: Configuring Interprovider Layer 3 VPN Option B	355
	Example: Configuring Interprovider Layer 3 VPN Option C	375
<b>Chapter 23</b>	<b>Configuring Carrier-of-Carrier VPNs</b>	<b>399</b>
	Understanding Carrier-of-Carriers VPNs	400
	Internet Service Provider as the Customer	401
	VPN Service Provider as the Customer	401
	Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service	402
	Configuring the Carrier-of-Carriers VPN Service Customer's CE Router	402
	Configuring MPLS	402
	Configuring BGP	403
	Configuring OSPF	403
	Configuring Policy Options	404
	Configuring the Carrier-of-Carriers VPN Service Provider's PE Routers	404
	Configuring MPLS	404
	Configuring BGP	405
	Configuring IS-IS	405
	Configuring LDP	405
	Configuring a Routing Instance	406

Configuring Policy Options . . . . .	406
Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service . . . . .	407
Configuring the Carrier-of-Carriers Customer's PE Router . . . . .	407
Configuring MPLS . . . . .	408
Configuring BGP . . . . .	408
Configuring OSPF . . . . .	408
Configuring LDP . . . . .	409
Configuring VPN Service in the Routing Instance . . . . .	409
Configuring Policy Options . . . . .	410
Configuring the Carrier-of-Carriers Customer's CE Router (or switch) . . . . .	410
Configuring MPLS . . . . .	410
Configuring BGP . . . . .	411
Configuring OSPF and LDP . . . . .	412
Configuring Policy Options . . . . .	412
Configuring the Provider's PE Router or Switch . . . . .	412
Configuring MPLS . . . . .	413
Configuring a PE-to-PE BGP Session . . . . .	413
Configuring IS-IS and LDP . . . . .	413
Configuring Policy Options . . . . .	414
Configuring a Routing Instance to Send Routes to the CE Router . . . . .	414
Carrier-of-Carriers VPN Example—Customer Provides Internet Service . . . . .	415
Network Topology for Carrier-of-Carriers Service . . . . .	415
Configuration for Router A . . . . .	416
Configuration for Router B . . . . .	416
Configuration for Router C . . . . .	417
Configuration for Router D . . . . .	418
Configuration for Router E . . . . .	419
Configuration for Router F . . . . .	420
Configuration for Router G . . . . .	420
Configuration for Router H . . . . .	421
Configuration for Router I . . . . .	422
Configuration for Router J . . . . .	423
Configuration for Router K . . . . .	423
Configuration for Router L . . . . .	424
Carrier-of-Carriers VPN Example—Customer Provides VPN Service . . . . .	425
Network Topology for Carrier-of-Carriers Service . . . . .	425
Configuration for Router A . . . . .	426
Configuration for Router B . . . . .	426
Configuration for Router C . . . . .	428
Configuration for Router D . . . . .	428
Configuration for Router E . . . . .	429
Configuration for Router F . . . . .	431
Configuration for Router G . . . . .	431
Configuration for Router H . . . . .	432
Configuration for Router I . . . . .	433
Configuration for Router J . . . . .	435
Configuration for Router K . . . . .	435

	Configuration for Router L . . . . .	436
	Multiple Instances for LDP and Carrier-of-Carriers VPNs . . . . .	436
<b>Part 4</b>	<b>Configuring Multicast on Layer 3 VPNs</b>	
<b>Chapter 24</b>	<b>Multicast VPN Overview . . . . .</b>	<b>441</b>
	Understanding MVPN Concepts and Protocols . . . . .	441
	Multicast over Layer 3 VPNs Overview . . . . .	441
	Sending PIM Hello Messages to the PE Routers . . . . .	443
	Sending PIM Join Messages to the PE Routers . . . . .	444
	Receiving the Multicast Transmission . . . . .	444
	MBGP Multicast VPN Sites . . . . .	445
	Example: Configuring MBGP Multicast VPNs . . . . .	446
	Configuring Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	464
	Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	465
	Configuring Selective Provider Tunnels for an MBGP MVPN . . . . .	466
	Configuring the Multicast Group Address for an MBGP MVPN . . . . .	468
	Configuring the Multicast Source Address for an MBGP MVPN . . . . .	468
	Configuring Static Selective Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	468
	Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	469
	Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	469
	Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	470
	Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview . . . . .	471
	Multicast VPN Terminology . . . . .	473
	Supported Multicast VPN Standards . . . . .	473
<b>Chapter 25</b>	<b>Configuration Overview . . . . .</b>	<b>475</b>
	Configuring Segmented Inter-Area P2MP LSP . . . . .	476
	Example: Configuring Segmented Inter-Area P2MP LSP . . . . .	478
<b>Chapter 26</b>	<b>Associating Routing Instances with Sender and Receiver Sites . . . . .</b>	<b>537</b>
	Configuring Routing Instances for an MBGP MVPN . . . . .	537
	Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs . . . . .	538
	Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs . . . . .	540
	Configuring Internet Multicast Using Ingress Replication Provider Tunnels . . . . .	542
	Example: Configuring PIM State Limits . . . . .	545
	Controlling PIM Resources for Multicast VPNs Overview . . . . .	546
	System Log Messages for PIM Resources . . . . .	547
	Example: Configuring PIM State Limits . . . . .	548
	Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN . . . . .	559
	About S-PMSI . . . . .	559
	Scenarios for Using Wildcard S-PMSI . . . . .	560
	Types of Wildcard S-PMSI . . . . .	561

	Differences Between Wildcard S-PMSI and (S,G) S-PMSI . . . . .	561
	Wildcard (*,*) S-PMSI and PIM Dense Mode . . . . .	561
	Wildcard (*,*) S-PMSI and PIM-BSR . . . . .	562
	Wildcard Source and the 0.0.0.0/0 Source Prefix . . . . .	562
	Configuring a Selective Provider Tunnel Using Wildcards . . . . .	564
	Example: Configuring Selective Provider Tunnels Using Wildcards . . . . .	565
<b>Chapter 27</b>	<b>Providing Resiliency with Redundant Virtual Tunnels . . . . .</b>	<b>567</b>
	Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic	
	Overview . . . . .	567
	Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering	
	Multicast Traffic . . . . .	569
	Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in	
	Delivering Multicast Traffic . . . . .	570
<b>Chapter 28</b>	<b>Creating Next Generation MVPN VRF Import and Export Policies . . . . .</b>	<b>587</b>
	Configuring VRF Route Targets for Routing Instances for an MBGP MVPN . . . . .	587
	Configuring the Export Target for an MBGP MVPN . . . . .	588
	Configuring the Import Target for an MBGP MVPN . . . . .	589
	Configuring the Import Target Receiver and Sender for an MBGP	
	MVPN . . . . .	589
	Configuring the Import Target Unicast Parameters for an MBGP	
	MVPN . . . . .	590
	Limiting Routes to Be Advertised by an MVPN VRF Instance . . . . .	590
<b>Chapter 29</b>	<b>Signaling Provider Tunnels in Next Generation MVPNs . . . . .</b>	<b>593</b>
	PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs . . . . .	593
	Configuring Nonstop Active Routing for BGP Multicast VPN . . . . .	594
	Configuring PIM Provider Tunnels for an MBGP MVPN . . . . .	596
	Configuring PIM-SSM GRE Selective Provider Tunnels . . . . .	597
	Configuring NLRI Parameters for an MBGP MVPN . . . . .	598
<b>Chapter 30</b>	<b>Associating Interfaces with MGBP VPNs . . . . .</b>	<b>599</b>
	Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs . . . . .	599
	Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs . . . . .	600
<b>Chapter 31</b>	<b>Configuring Draft Rosen VPNs . . . . .</b>	<b>611</b>
	Configuring Multicast Layer 3 VPNs . . . . .	611
	Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast	
	VPN . . . . .	612
<b>Chapter 32</b>	<b>Configuring Load Balancing . . . . .</b>	<b>623</b>
	Example: Configuring PIM Join Load Balancing on Next-Generation Multicast	
	VPN . . . . .	623
<b>Chapter 33</b>	<b>Tracing Traffic and Operations . . . . .</b>	<b>633</b>
	Tracing MBGP MVPN Traffic and Operations . . . . .	633



<b>Part 5</b>	<b>Configuring Full-Mesh, Hub-and-Spoke, and Overlapping VPNs</b>	
<b>Chapter 34</b>	<b>Configuring Full-Mesh VPNs</b>	<b>637</b>
	Configuring a Simple Full-Mesh VPN Topology	637
	Enabling an IGP on the PE and P Routers	638
	Enabling RSVP and MPLS on the P Router	639
	Configuring the MPLS LSP Tunnel Between the PE Routers	639
	Configuring IBGP on the PE Routers	640
	Configuring Routing Instances for VPNs on the PE Routers	641
	Configuring VPN Policy on the PE Routers	643
	Simple VPN Configuration Summarized by Router	646
	Router A (PE Router)	646
	Router B (P Router)	649
	Router C (PE Router)	649
	Configuring a Full-Mesh VPN Topology with Route Reflectors	651
<b>Chapter 35</b>	<b>Configuring Hub-and-Spoke VPNs</b>	<b>653</b>
	Configuring Hub-and-Spoke VPN Topologies: One Interface	653
	Configuring Hub CE1	654
	Configuring Hub PE1	655
	Configuring the P Router	655
	Configuring Spoke PE2	656
	Configuring Spoke PE3	658
	Configuring Spoke CE2	659
	Configuring Spoke CE3	659
	Enabling Egress Features on the Hub PE Router	661
	Configuring Hub PE1	662
	Configuring Hub-and-Spoke VPN Topologies: Two Interfaces	665
	Enabling an IGP on the Hub-and-Spoke PE Routers	667
	Configuring LDP on the Hub-and-Spoke PE Routers	668
	Configuring IBGP on the PE Routers	668
	Configuring VPN Routing Instances on the Hub-and-Spoke PE Routers	669
	Configuring VPN Policy on the PE Routers	672
	Hub-and-Spoke VPN Configuration Summarized by Router	675
	Router D (Hub PE Router)	675
	Router E (Spoke PE Router)	677
	Router F (Spoke PE Router)	678
<b>Chapter 36</b>	<b>Configuring Overlapping VPNs</b>	<b>681</b>
	Configuring Overlapping VPNs Using Routing Table Groups	681
	Configuring Routing Table Groups	682
	Configuring Static Routes Between the PE and CE Routers	683
	Configuring the Routing Instance for VPN A	683
	Configuring the Routing Instance for VPN AB	684
	Configuring the Routing Instance for VPN B	684
	Configuring VPN Policy	685
	Configuring BGP Between the PE and CE Routers	688
	Configuring OSPF Between the PE and CE Routers	689

	Configuring Static, BGP, and OSPF Routes Between PE and CE Routers . . .	691
	Configuring Overlapping VPNs Using Automatic Route Export . . . . .	692
	Configuring Overlapping VPNs with BGP and Automatic Route Export . . .	693
	Configuring Overlapping VPNs and Additional Tables . . . . .	694
	Configuring Automatic Route Export for All VRF Instances . . . . .	695
<b>Part 6</b>	<b>Configuring IPsec VPNs in Layer 3 VPNs</b>	
<b>Chapter 37</b>	<b>Configuring IPsec Tunnels . . . . .</b>	<b>699</b>
	Configuring IPsec Tunnels Instead of MPLS LSPs Between PE Routers in Layer 3 VPNs . . . . .	699
<b>Chapter 38</b>	<b>Configuring ES Tunnel Interfaces for Layer 3 VPNs . . . . .</b>	<b>703</b>
	Configuring an ES Tunnel Interface for Layer 3 VPNs . . . . .	703
	Configuring the ES Tunnel Interface on the PE Router . . . . .	703
	Configuring the ES Tunnel Interface on the CE Router . . . . .	704
	Configuring an ES Tunnel Interface Between a PE and CE Router . . . . .	705
	Configuring IPsec on Router PE1 . . . . .	705
	Configuring the Routing Instance Without the Encapsulating Interface . .	706
	Configuring the Routing Instance on Router PE1 . . . . .	706
	Configuring the ES Tunnel Interface on Router PE1 . . . . .	706
	Configuring the Encapsulating Interface for the ES Tunnel . . . . .	707
	Configuring the Routing Instance with the Encapsulating Interface . . . . .	707
	Configuring the Routing Instance on Router PE1 . . . . .	707
	Configuring the ES Tunnel Interface on Router PE1 . . . . .	708
	Configuring the Encapsulating Interface on Router PE1 . . . . .	708
	Configuring the ES Tunnel Interface on Router CE1 . . . . .	708
	Configuring IPsec on Router CE1 . . . . .	708
<b>Chapter 39</b>	<b>Configuring GRE Tunnel Interfaces for Layer 3 VPNs . . . . .</b>	<b>711</b>
	Configuring GRE Tunnels for Layer 3 VPNs . . . . .	711
	Configuring GRE Tunnels Manually Between PE and CE Routers . . . . .	712
	Configuring the GRE Tunnel Interface on the PE Router . . . . .	712
	Configuring the GRE Tunnel Interface on the CE Router . . . . .	713
	Configuring GRE Tunnels Dynamically . . . . .	713
	Configuring a GRE Tunnel Interface Between PE Routers . . . . .	714
	Configuring the Routing Instance on Router A . . . . .	715
	Configuring the Routing Instance on Router D . . . . .	716
	Configuring MPLS, BGP, and OSPF on Router A . . . . .	716
	Configuring MPLS, BGP, and OSPF on Router D . . . . .	716
	Configuring the Tunnel Interface on Router A . . . . .	717
	Configuring the Tunnel Interface on Router D . . . . .	717
	Configuring the Routing Options on Router A . . . . .	718
	Configuring the Routing Options on Router D . . . . .	718
	Configuration Summary for Router A . . . . .	718
	Configuration Summary for Router D . . . . .	720
	Configuring a GRE Tunnel Interface Between a PE and CE Router . . . . .	721
	Configuring the Routing Instance Without the Encapsulating Interface . . .	722
	Configuring the Routing Instance on Router PE1 . . . . .	722
	Configuring the GRE Tunnel Interface on Router PE1 . . . . .	722

	Configuring the Encapsulation Interface on Router PE1 . . . . .	723
	Configuring the Routing Instance with the Encapsulating Interface . . . . .	723
	Configuring the Routing Instance on Router PE1 . . . . .	723
	Configuring the GRE Tunnel Interface on Router PE1 . . . . .	724
	Configuring the Encapsulation Interface on Router PE1 . . . . .	725
	Configuring the GRE Tunnel Interface on Router CE1 . . . . .	725
<b>Chapter 40</b>	<b>Configuring Next-Hop-Based Dynamic Tunnels for Layer 3 VPNs . . . . .</b>	<b>727</b>
	Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels . . . . .	727
	Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels . . .	740
	Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview . . .	753
	Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels . . . . .	756
<b>Part 7</b>	<b>Configuring Load Balancing in Layer 3 VPNs</b>	
<b>Chapter 41</b>	<b>Configuring Load Balancing in Layer 3 VPNs . . . . .</b>	<b>769</b>
	VPN Per-Packet Load Balancing . . . . .	769
	Load Balancing and IP Header Filtering for Layer 3 VPNs . . . . .	770
	Layer 3 VPN Load Balancing Using IP Header Filtering . . . . .	771
	Layer 3 VPN Load Balancing Overview . . . . .	771
	Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering . . . . .	772
	Configuring Protocol-Independent Load Balancing in Layer 3 VPNs . . . . .	787
	Configuring Load Balancing for Layer 3 VPNs . . . . .	787
	Configuring Load Balancing and Routing Policies . . . . .	789
<b>Chapter 42</b>	<b>Enabling BGP PIC for Layer 3 VPNs . . . . .</b>	<b>791</b>
	Configuring BGP PIC Edge for MPLS Layer 3 VPNs . . . . .	791
	Example: Configuring BGP PIC Edge for MPLS Layer 3 VPNs . . . . .	793
<b>Part 8</b>	<b>Configuring Protection and Performance Features for Layer 3 VPNs</b>	
<b>Chapter 43</b>	<b>Configuring Egress Protection in Layer 3 VPNs . . . . .</b>	<b>809</b>
	Egress Protection for BGP Labeled Unicast . . . . .	809
	Configuring Egress Protection for BGP Labeled Unicast . . . . .	811
	Example: Configuring Egress Protection for BGP Labeled Unicast . . . . .	813
	Example: Configuring MPLS Egress Protection for Layer 3 VPN Services . . . . .	826
	Egress Protection for Layer 3 VPN Edge Protection Overview . . . . .	827
	Router Functions . . . . .	828
	Protector and Protection Models . . . . .	829
	IGP Advertisement Model . . . . .	830
	Example: Configuring Egress Protection for Layer 3 VPN Services . . . . .	834
	Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP . . . . .	843

<b>Chapter 44</b>	<b>Configuring Provider Edge Link Protection in Layer 3 VPNs . . . . . 877</b>
	Understanding Provider Edge Link Protection for BGP Labeled Unicast Paths . . 877
	Example: Configuring Provider Edge Link Protection in Layer 3 VPNs . . . . . 878
	Understanding Provider Edge Link Protection in Layer 3 VPNs . . . . . 879
	Example: Configuring Provider Edge Link Protection in Layer 3 VPNs . . . . . 880
	Example: Configuring Provider Edge Link Protection for BGP Labeled Unicast Paths . . . . . 895
<b>Chapter 45</b>	<b>Improving Layer 3 VPN Performance . . . . . 911</b>
	Example: Configuring Link Protection with Host Fast Reroute . . . . . 911
	Understanding Host Fast Reroute . . . . . 911
	ARP Prefix Limit and Blackout Supplementary Timeout . . . . . 913
	Primary Route and Backup Route Candidates . . . . . 914
	Backup Path Selection Policy . . . . . 914
	Characteristics of HFRR Routes . . . . . 915
	Removal of HFRR Routes . . . . . 915
	Interfaces That Support HFRR . . . . . 915
	Example: Configuring Link Protection with Host Fast Reroute . . . . . 916
	Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs . . . . 928
	Accepting Up to One Million Layer 3 VPN Route Updates . . . . . 929
	Accepting More Than One Million Layer 3 VPN Route Updates . . . . . 930
	Enabling Chained Composite Next Hops for IPv6 Labeled Unicast Routes . . . . . 932
<b>Part 9</b>	<b>Connecting Layer 3 VPNs to Other VPNs</b>
<b>Chapter 46</b>	<b>Connecting Layer 3 VPNs to Layer 2 Circuits . . . . . 935</b>
	Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN . . . . . 935
	Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN . . . . . 936
<b>Chapter 47</b>	<b>Connecting Layer 3 VPNs to Layer 2 VPNs . . . . . 957</b>
	Interconnecting Layer 2 VPNs with Layer 3 VPNs Overview . . . . . 957
	Interconnecting Layer 2 VPNs with Layer 3 VPNs Applications . . . . . 957
	Example: Interconnecting a Layer 2 VPN with a Layer 3 VPN . . . . . 958
<b>Part 10</b>	<b>Troubleshooting Layer 3 VPNs</b>
<b>Chapter 48</b>	<b>Troubleshooting Layer 3 VPNs . . . . . 985</b>
	Diagnosing Common Layer 3 VPN Problems . . . . . 985
	Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing . . . . . 989
	Example: Troubleshooting Layer 3 VPNs . . . . . 996
<b>Part 11</b>	<b>Configuration Statements and Operational Commands</b>
<b>Chapter 49</b>	<b>Configuration Statements (All VPNs) . . . . . 1009</b>
	aggregate-label . . . . . 1010
	backup-neighbor . . . . . 1011
	description (Routing Instances) . . . . . 1012
	family route-target . . . . . 1013

	graceful-restart (Enabling Globally) .....	1015
	instance-type .....	1017
	interface (Routing Instances) .....	1019
	no-forwarding .....	1020
	proxy-generate .....	1021
	revert-time (Protocols Layer 2 Circuits) .....	1022
	route-distinguisher .....	1024
	route-distinguisher-id .....	1026
	route-target-filter .....	1027
	switchover-delay .....	1028
	unicast-reverse-path .....	1029
	vpn-apply-export .....	1030
	vrf-export .....	1031
	vrf-import .....	1032
	vrf-mtu-check .....	1033
	vrf-target .....	1034
<b>Chapter 50</b>	<b>Configuration Statements (Layer 3 VPNs) .....</b>	<b>1037</b>
	advertise-from-main-vpn-tables .....	1040
	advertise-mode (MPLS) .....	1041
	all-regions .....	1042
	arp-prefix-limit (Host Fast Reroute) .....	1043
	as-path-compare .....	1044
	chained-composite-next-hop .....	1045
	classifiers .....	1047
	create-new-ucast-tunnel .....	1048
	domain-id .....	1049
	domain-vpn-tag .....	1050
	dynamic-tunnels .....	1051
	egress-protection (BGP) .....	1052
	egress-protection (MPLS) .....	1053
	egress-protection (Routing Instances) .....	1054
	export-target .....	1054
	extended-space .....	1055
	family (VRF Advertisement) .....	1056
	forwarding-table .....	1057
	forwarding-context (Protocols BGP) .....	1058
	global-arp-prefix-limit (Host Fast Reroute) .....	1059
	global-supplementary-blackout-timer (Host Fast Reroute) .....	1061
	group (Routing Instances) .....	1063
	group-address (Routing Instances VPN) .....	1065
	group-range (MBGP MVPN Tunnel) .....	1066
	group-rp-mapping .....	1067
	host-fast-reroute .....	1068
	import-target .....	1069
	independent-domain .....	1070
	inet-mvpn (BGP) .....	1071
	inet-mvpn (VRF Advertisement) .....	1072
	inet6 .....	1072

inet6-mvpn (BGP) .....	1073
inet6-mvpn (VRF Advertisement) .....	1074
inet6-vpn .....	1075
ingress (Chained Composite Next Hop) .....	1076
ingress-replication .....	1078
interface (Host Fast Reroute) .....	1079
ip-tunnel-rpf-check .....	1080
interface (Virtual Tunnel in Routing Instances) .....	1081
inter-region .....	1082
inter-region-segmented .....	1083
inter-region-template .....	1085
l3vpn .....	1087
l3vpn (transit) .....	1088
label .....	1089
label-switched-path (Protocols MVPN) .....	1090
label-switched-path-template (Multicast) .....	1091
labeled-bgp .....	1092
labeled-unicast .....	1093
link-protection (Host Fast Reroute) .....	1094
localized-fib .....	1095
maximum-paths .....	1096
maximum-prefixes .....	1098
metric (Protocols OSPF Sham Link) .....	1099
mpls-internet-multicast .....	1100
multicast (Virtual Tunnel in Routing Instances) .....	1100
multihop .....	1101
multipath (Routing Options) .....	1103
mvpn .....	1104
mvpn-mode .....	1106
no-vrf-advertise .....	1107
no-vrf-propagate-ttl .....	1108
per-group-label .....	1109
pim-asm .....	1109
pim-ssm (Selective Tunnel) .....	1110
primary (Virtual Tunnel in Routing Instances) .....	1111
protect core .....	1112
protection (Protocols BGP) .....	1113
provider-tunnel .....	1114
redundancy-group (Chassis - MX Series) .....	1118
redundancy-group (Interfaces) .....	1119
redundant-logical-tunnel .....	1120
redundant-virtual-tunnel .....	1121
region .....	1122
register-limit .....	1123
route-target (Protocols MVPN) .....	1124
routing-instances (CoS) .....	1125
rpt-spt .....	1126
rsvp-te (Routing Instances Provider Tunnel Selective) .....	1127
rsvp-te (Protocols MVPN) .....	1128

selective	1129
sglimit	1131
sham-link	1132
sham-link-remote	1133
source (Routing Instances Provider Tunnel Selective)	1134
source-class-usage	1135
spt-only	1136
static-lsp	1137
supplementary-blackout-timer (Host Fast Reroute)	1139
target (Routing Instances MVPN)	1140
template(Protocols MVPN)	1141
threshold-rate	1143
traceoptions (Protocols MVPN)	1144
traffic-statistics (Protocols BGP)	1146
tunnel-limit (Routing Instances Provider Tunnel Selective)	1147
unicast (Route Target Community)	1148
unicast (Virtual Tunnel in Routing Instances)	1149
vpn-localization	1150
vpn-unequal-cost	1151
vrf-advertise-selective	1152
vrf-propagate-ttl	1153
vrf-table-label	1154
wildcard-group-inet	1156
wildcard-group-inet6	1157
wildcard-source (Selective Provider Tunnels)	1158
<b>Chapter 51</b>	
<b>Operational Commands</b>	<b>1159</b>
ping mpls l3vpn	1160
show hfr profiles	1163
show ingress-replication mvpn	1165
show mvpn c-multicast	1167
show mvpn instance	1170
show mvpn neighbor	1174
show route vpn-localization	1179





# List of Figures

<b>Part 1</b>	<b>Common Configuration for All VPNs</b>	
<b>Chapter 1</b>	<b>VPNs Overview</b>	<b>3</b>
	Figure 1: Logical Interface per Router in a Virtual-Router Routing Instance	6
	Figure 2: Routers in a VPN	8
<b>Chapter 3</b>	<b>Distributing Routes in VPNs</b>	<b>19</b>
	Figure 3: Network Topology of Site of Origin Example	32
<b>Chapter 4</b>	<b>Distributing VPN Routes with Target Filtering</b>	<b>37</b>
	Figure 4: BGP Route Target Filtering Enabled for a Group of VPNs	41
	Figure 5: Proxy BGP Route Target Filtering Topology	53
	Figure 6: BGP Route Target Filtering Export Policy Topology	69
<b>Chapter 5</b>	<b>Configuring Forwarding Options for VPNs</b>	<b>85</b>
	Figure 7: Chained Composite Next Hop for PE-PE Connections	88
<b>Chapter 6</b>	<b>Configuring Graceful Restart for VPNs</b>	<b>101</b>
	Figure 8: Unicast RPF Sample Topology	107
<b>Part 2</b>	<b>Common Configuration for Layer 3 VPNs</b>	
<b>Chapter 9</b>	<b>Layer 3 VPN Overview</b>	<b>123</b>
	Figure 9: General Layer 3 VPN Architecture	126
	Figure 10: Using MPLS LSPs to Tunnel Between PE Routers	128
	Figure 11: Label Stack	128
	Figure 12: VPN Attributes and Route Distribution	129
<b>Chapter 11</b>	<b>Creating Unique VPN Routes Using VRF Tables</b>	<b>141</b>
	Figure 13: VRF Tables	142
	Figure 14: Example VRF Localization	148
<b>Chapter 12</b>	<b>Distributing Routes in Layer 3 VPNs</b>	<b>169</b>
	Figure 15: OSPF Sham Link	172
	Figure 16: Example of a Configuration Using an OSPF Domain ID	180
	Figure 17: OSPFv2 Sham Link	185
	Figure 18: OSPFv2 Sham Link Example	187
	Figure 19: Example of an LDP-over-RSVP VPN Topology	196
	Figure 20: Label Pushing and Popping	198
<b>Chapter 13</b>	<b>Associating Interfaces with Layer 3 VPNs</b>	<b>211</b>
	Figure 21: Application-Based Layer 3 VPN Example Configuration	212
<b>Chapter 14</b>	<b>Carrying IPv4 Traffic Over Layer 3 VPNs</b>	<b>217</b>

	Figure 22: Overlapping Addresses Among Different VPNs . . . . .	218
	Figure 23: Route Distinguishers . . . . .	220
	Figure 24: Route Distribution Within a VPN . . . . .	221
	Figure 25: Distribution of Routes from CE Routers to PE Routers . . . . .	222
	Figure 26: Distribution of Routes Between PE Routers . . . . .	223
	Figure 27: Distribution of Routes from PE Routers to CE Routers . . . . .	225
<b>Chapter 15</b>	<b>Carrying IPv6 Traffic Over Layer 3 VPNs . . . . .</b>	<b>227</b>
	Figure 28: Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains . . . . .	233
<b>Chapter 16</b>	<b>Configuring an AS for Layer 3 VPNs . . . . .</b>	<b>241</b>
	Figure 29: AS Override Topology . . . . .	243
<b>Chapter 20</b>	<b>Enabling Internet Access for Layer 3 VPNs . . . . .</b>	<b>271</b>
	Figure 30: PE Router Does Not Provide Internet Access . . . . .	272
	Figure 31: PE Router Connects to a Router Connected to the Internet . . . . .	272
	Figure 32: Routing VPN and Internet Traffic Through Different Interfaces . . . . .	273
	Figure 33: Example of Internet Traffic Routed Through Separate Interfaces . . . . .	273
	Figure 34: VPN and Outgoing Internet Traffic Routed Through the Same Interface and Return Internet Traffic Routed Through a Different Interface . . . . .	279
	Figure 35: Interface Configured to Carry Both Internet and VPN Traffic . . . . .	280
	Figure 36: VPN and Internet Traffic Routed Through the Same Interface . . . . .	284
	Figure 37: Internet Traffic Routed Through a Separate NAT Device . . . . .	288
	Figure 38: Internet Traffic Routed Through a NAT Example Topology . . . . .	289
	Figure 39: Internet Access Through a Hub CE Router Performing NAT . . . . .	297
	Figure 40: Internet Access Provided Through a Hub CE Router . . . . .	297
	Figure 41: Two Hub CE Routers Handling Internet Traffic and NAT . . . . .	301
<b>Part 3</b>	<b>Configuring Interprovider and Carrier-of-Carrier VPNs</b>	
<b>Chapter 22</b>	<b>Configuring Interprovider VPNs . . . . .</b>	<b>317</b>
	Figure 42: Interprovider VPN Network Topology . . . . .	317
	Figure 43: Network Topology for the Interprovider VPN Example . . . . .	320
	Figure 44: Network Topology of Interprovider VPN Example—Multihop MP-EBGP . . . . .	327
	Figure 45: Physical Topology of Interprovider Layer 3 VPN Option A . . . . .	335
	Figure 46: Physical Topology of Interprovider Layer 3 VPN Option B . . . . .	356
	Figure 47: Physical Topology of Interprovider Layer 3 VPN Option C . . . . .	377
<b>Chapter 23</b>	<b>Configuring Carrier-of-Carrier VPNs . . . . .</b>	<b>399</b>
	Figure 48: Carrier-of-Carriers VPN Architecture . . . . .	400
	Figure 49: Carrier-of-Carriers VPN Example Network Topology . . . . .	416
	Figure 50: Carrier-of-Carriers VPN Example Network Topology . . . . .	426
<b>Part 4</b>	<b>Configuring Multicast on Layer 3 VPNs</b>	
<b>Chapter 24</b>	<b>Multicast VPN Overview . . . . .</b>	<b>441</b>
	Figure 51: Multicast Topology Overview . . . . .	443
	Figure 52: Multicast Over Layer 3 VPN Example Topology . . . . .	447
<b>Chapter 25</b>	<b>Configuration Overview . . . . .</b>	<b>475</b>

	Figure 53: Example Segmented Inter-Area P2MP LSP . . . . .	479
<b>Chapter 26</b>	<b>Associating Routing Instances with Sender and Receiver Sites . . . . .</b>	<b>537</b>
	Figure 54: Internet Multicast Topology . . . . .	543
	Figure 55: PIM State Limits Topology . . . . .	549
	Figure 56: Simple MVPN Topology . . . . .	560
<b>Chapter 27</b>	<b>Providing Resiliency with Redundant Virtual Tunnels . . . . .</b>	<b>567</b>
	Figure 57: Redundant VT Interfaces in MBGP MVPN . . . . .	572
<b>Chapter 30</b>	<b>Associating Interfaces with MGBP VPNs . . . . .</b>	<b>599</b>
	Figure 58: Multiple VT Interfaces in MBGP MVPN Topology . . . . .	600
<b>Chapter 31</b>	<b>Configuring Draft Rosen VPNs . . . . .</b>	<b>611</b>
	Figure 59: PIM Join Load Balancing on Draft-Rosen MVPN . . . . .	617
<b>Chapter 32</b>	<b>Configuring Load Balancing . . . . .</b>	<b>623</b>
	Figure 60: PIM Join Load Balancing on Next-Generation MVPN . . . . .	626
<b>Part 5</b>	<b>Configuring Full-Mesh, Hub-and-Spoke, and Overlapping VPNs</b>	
<b>Chapter 34</b>	<b>Configuring Full-Mesh VPNs . . . . .</b>	<b>637</b>
	Figure 61: Example of a Simple VPN Topology . . . . .	637
<b>Chapter 35</b>	<b>Configuring Hub-and-Spoke VPNs . . . . .</b>	<b>653</b>
	Figure 62: Example of a Hub-and-Spoke VPN Topology with One Interface . . .	653
	Figure 63: Example of a Hub-and-Spoke VPN Topology with Two Interfaces . .	666
	Figure 64: Route Distribution Between Two Spoke Routers . . . . .	667
<b>Chapter 36</b>	<b>Configuring Overlapping VPNs . . . . .</b>	<b>681</b>
	Figure 65: Example of an Overlapping VPN Topology . . . . .	682
<b>Part 6</b>	<b>Configuring IPsec VPNs in Layer 3 VPNs</b>	
<b>Chapter 38</b>	<b>Configuring ES Tunnel Interfaces for Layer 3 VPNs . . . . .</b>	<b>703</b>
	Figure 66: ES Tunnel Interface (IPsec Tunnel) . . . . .	705
<b>Chapter 39</b>	<b>Configuring GRE Tunnel Interfaces for Layer 3 VPNs . . . . .</b>	<b>711</b>
	Figure 67: GRE Tunnel Configured Between the Local CE Router and the PE Router . . . . .	711
	Figure 68: GRE Tunnel Configured Between the Remote CE Router and the PE Router . . . . .	711
	Figure 69: PE Routers A and D Connected by a GRE Tunnel Interface . . . . .	715
	Figure 70: GRE Tunnel Between the CE Router and the PE Router . . . . .	721
<b>Chapter 40</b>	<b>Configuring Next-Hop-Based Dynamic Tunnels for Layer 3 VPNs . . . . .</b>	<b>727</b>
	Figure 71: Layer 3 VPN over Dynamic GRE Tunnels . . . . .	729
	Figure 72: Dynamic MPLS-over-UDP Tunnels . . . . .	742
	Figure 73: Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels . . .	754
	Figure 74: Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels . . .	757
<b>Part 7</b>	<b>Configuring Load Balancing in Layer 3 VPNs</b>	
<b>Chapter 41</b>	<b>Configuring Load Balancing in Layer 3 VPNs . . . . .</b>	<b>769</b>

	Figure 75: Layer 3 VPN Load Balancing Using IP Header Filtering . . . . .	773
<b>Chapter 42</b>	<b>Enabling BGP PIC for Layer 3 VPNs . . . . .</b>	<b>791</b>
	Figure 76: BGP PIC Edge Scenario . . . . .	795
<b>Part 8</b>	<b>Configuring Protection and Performance Features for Layer 3 VPNs</b>	
<b>Chapter 43</b>	<b>Configuring Egress Protection in Layer 3 VPNs . . . . .</b>	<b>809</b>
	Figure 77: Inter-AS Option C . . . . .	810
	Figure 78: Seamless MPLS . . . . .	810
	Figure 79: Egress Protection in a Layer 3 VPN . . . . .	814
	Figure 80: Sample Topology for Egress Protection . . . . .	827
	Figure 81: Layer 3 VPN Egress Protection with RSVP and LDP . . . . .	844
<b>Chapter 44</b>	<b>Configuring Provider Edge Link Protection in Layer 3 VPNs . . . . .</b>	<b>877</b>
	Figure 82: Provider Edge Link Protection in a Layer 3 VPN . . . . .	881
	Figure 83: Labeled Unicast Link Protection in a Layer 3 VPN . . . . .	896
<b>Chapter 45</b>	<b>Improving Layer 3 VPN Performance . . . . .</b>	<b>911</b>
	Figure 84: Host Fast Reroute . . . . .	912
	Figure 85: Host Fast Reroute Topology . . . . .	918
<b>Part 9</b>	<b>Connecting Layer 3 VPNs to Other VPNs</b>	
<b>Chapter 46</b>	<b>Connecting Layer 3 VPNs to Layer 2 Circuits . . . . .</b>	<b>935</b>
	Figure 86: Physical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection . . . . .	936
	Figure 87: Logical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection . . . . .	937
<b>Chapter 47</b>	<b>Connecting Layer 3 VPNs to Layer 2 VPNs . . . . .</b>	<b>957</b>
	Figure 88: Physical Topology of a Layer 2 VPN Terminating into a Layer 3 VPN . . . . .	960
	Figure 89: Logical Topology of a Layer 2 VPN Terminating into a Layer 3 VPN . .	960
<b>Part 10</b>	<b>Troubleshooting Layer 3 VPNs</b>	
<b>Chapter 48</b>	<b>Troubleshooting Layer 3 VPNs . . . . .</b>	<b>985</b>
	Figure 90: Disabling TTL Propagation for a Single VPN . . . . .	990
	Figure 91: Layer 3 VPN Topology for ping and traceroute Examples . . . . .	997

# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xxvii</b>
	Table 1: Notice Icons . . . . .	xxix
	Table 2: Text and Syntax Conventions . . . . .	xxx
<b>Part 1</b>	<b>Common Configuration for All VPNs</b>	
<b>Chapter 5</b>	<b>Configuring Forwarding Options for VPNs</b> . . . . .	<b>85</b>
	Table 3: Support for Chained Composite Next Hops . . . . .	85
<b>Part 2</b>	<b>Common Configuration for Layer 3 VPNs</b>	
<b>Chapter 11</b>	<b>Creating Unique VPN Routes Using VRF Tables</b> . . . . .	<b>141</b>
	Table 4: Support for Aggregated and VLAN Interfaces . . . . .	161
	Table 5: Support for ATM and Frame Relay Interfaces . . . . .	161
	Table 6: Support for Ethernet, SONET/SDH, and T1/T3/E3 Interfaces . . . . .	162
	Table 7: Support for Channelized IQE Interfaces on M320 Routers with Enhanced III FPCs . . . . .	163
	Table 8: Support for Multilink PPP and Multilink Frame Relay Interfaces . . . . .	164
<b>Chapter 12</b>	<b>Distributing Routes in Layer 3 VPNs</b> . . . . .	<b>169</b>
	Table 9: How a PE Router Redistributes and Advertises Routes . . . . .	174
<b>Part 3</b>	<b>Configuring Interprovider and Carrier-of-Carrier VPNs</b>	
<b>Chapter 23</b>	<b>Configuring Carrier-of-Carrier VPNs</b> . . . . .	<b>399</b>
	Table 10: Comparison of Interprovider and Carrier-of-Carriers VPNs . . . . .	401
<b>Part 4</b>	<b>Configuring Multicast on Layer 3 VPNs</b>	
<b>Chapter 26</b>	<b>Associating Routing Instances with Sender and Receiver Sites</b> . . . . .	<b>537</b>
	Table 11: PIM System Log Messages . . . . .	548
<b>Part 7</b>	<b>Configuring Load Balancing in Layer 3 VPNs</b>	
<b>Chapter 41</b>	<b>Configuring Load Balancing in Layer 3 VPNs</b> . . . . .	<b>769</b>
	Table 12: Ingress Router Hashing . . . . .	770
	Table 13: Transit and Egress Router Hashing . . . . .	770
	Table 14: Device IP Address Quick Reference . . . . .	774
<b>Part 11</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 51</b>	<b>Operational Commands</b> . . . . .	<b>1159</b>

Table 15: show hfr profiles Output Fields . . . . .	1163
Table 16: show ingress-replication mvpn Output Fields . . . . .	1165
Table 17: show mvpn c-multicast Output Fields . . . . .	1167
Table 18: show mvpn instance Output Fields . . . . .	1170
Table 19: show mvpn neighbor Output Fields . . . . .	1174
Table 20: show route vpn-localization Output Fields . . . . .	1179

# About the Documentation

- [Documentation and Release Notes on page xxvii](#)
- [Supported Platforms on page xxvii](#)
- [Using the Examples in This Manual on page xxvii](#)
- [Documentation Conventions on page xxix](#)
- [Documentation Feedback on page xxxi](#)
- [Requesting Technical Support on page xxxi](#)

## 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 <https://www.juniper.net/documentation/>.

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 <https://www.juniper.net/books>.

## Supported Platforms

---

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

- [ACX Series](#)
- [M Series](#)
- [MX Series](#)
- [PTX Series](#)
- [T Series](#)

## 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
```

## 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.xml; }
```

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 xxix defines notice icons used in this guide.

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 xxx defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast</b>   <b>multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp</b> { <b>#</b> Required for dynamic MPLS only
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members</b> [ <i>community-ids</i> ]
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
<b>GUI Conventions</b>		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## 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 <https://www.juniper.net/documentation/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 <https://www.juniper.net/documentation/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto: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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://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 <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

## PART 1

# Common Configuration for All VPNs

- [VPNs Overview on page 3](#)
- [Assigning Routing Instances to VPNs on page 9](#)
- [Distributing Routes in VPNs on page 19](#)
- [Distributing VPN Routes with Target Filtering on page 37](#)
- [Configuring Forwarding Options for VPNs on page 85](#)
- [Configuring Graceful Restart for VPNs on page 101](#)
- [Configuring Class of Service for VPNs on page 115](#)
- [Pinging VPNs on page 117](#)



## CHAPTER 1

# VPNs Overview

- [VPLS on page 3](#)
- [Types of VPNs on page 4](#)
- [VPNs and Logical Systems on page 7](#)
- [Layer 2 VPNs on page 7](#)
- [Routers in a VPN on page 8](#)

## VPLS

---

In a Layer 3 network only, you can configure virtual private LAN service (VPLS), which is an Ethernet-based point-to-multipoint Layer 2 VPN. It enables you to connect geographically dispersed Ethernet local area networks (LAN) sites to each other across an MPLS backbone. For ISP customers who implement VPLS, all sites appear to be in the same Ethernet LAN even though traffic travels across the service provider's network.

- Related Documentation**
- *Junos OS VPNs Library for Routing Devices*
  - *MX Series Router Architecture*

## Types of VPNs

---

A virtual private network (VPN) consists of two topological areas: the provider's network and the customer's network. The customer's network is commonly located at multiple physical sites and is also private (non-Internet). A customer site would typically consist of a group of routers or other networking equipment located at a single physical location. The provider's network, which runs across the public Internet infrastructure, consists of routers that provide VPN services to a customer's network as well as routers that provide other services. The provider's network connects the various customer sites in what appears to the customer and the provider to be a private network.

To ensure that VPNs remain private and isolated from other VPNs and from the public Internet, the provider's network maintains policies that keep routing information from different VPNs separate. A provider can service multiple VPNs as long as its policies keep routes from different VPNs separate. Similarly, a customer site can belong to multiple VPNs as long as it keeps routes from the different VPNs separate.

The Junos<sup>®</sup> Operating System (Junos OS) provides several types of VPNs; you can choose the best solution for your network environment. Each of the following VPNs has different capabilities and requires different types of configuration:

- [Layer 2 VPNs on page 4](#)
- [Layer 3 VPNs on page 5](#)
- [VPLS on page 5](#)
- [Virtual-Router Routing Instances on page 6](#)

### Layer 2 VPNs

Implementing a Layer 2 VPN on a router is similar to implementing a VPN using a Layer 2 technology such as ATM or Frame Relay. However, for a Layer 2 VPN on a router, traffic is forwarded to the router in Layer 2 format. It is carried by MPLS over the service provider's network and then converted back to Layer 2 format at the receiving site. You can configure different Layer 2 formats at the sending and receiving sites. The security and privacy of an MPLS Layer 2 VPN are equal to those of an ATM or Frame Relay VPN.

On a Layer 2 VPN, routing occurs on the customer's routers, typically on the CE router. The CE router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The PE router receiving the traffic sends it across the service provider's network to the PE router connected to the receiving site. The PE routers do not need to store or process the customer's routes; they only need to be configured to send data to the appropriate tunnel.

For a Layer 2 VPN, customers need to configure their own routers to carry all Layer 3 traffic. The service provider needs to know only how much traffic the Layer 2 VPN needs to carry. The service provider's routers carry traffic between the customer's sites using Layer 2 VPN interfaces. The VPN topology is determined by policies configured on the PE routers.



## Layer 3 VPNs

In a Layer 3 VPN, the routing occurs on the service provider's routers. Therefore, Layer 3 VPNs require more configuration on the part of the service provider, because the service provider's PE routers must store and process the customer's routes.

In the Junos OS, Layer 3 VPNs are based on RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*. This RFC defines a mechanism by which service providers can use their IP backbones to provide Layer 3 VPN services to their customers. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

VPNs based on RFC 4364 are also known as BGP/MPLS VPNs because BGP is used to distribute VPN routing information across the provider's backbone, and MPLS is used to forward VPN traffic across the backbone to remote VPN sites.

Customer networks, because they are private, can use either public addresses or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the private addresses used by other network users. BGP/MPLS VPNs solve this problem by prefixing a VPN identifier to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and within the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only.

## VPLS

Virtual private LAN service (VPLS) allows you to connect geographically dispersed customer sites as if they were connected to the same LAN. In many ways, it works like a Layer 2 VPN. VPLS and Layer 2 VPNs use the same network topology and function similarly. A packet originating within a customer's network is sent first to a CE device. It is then sent to a PE router within the service provider's network. The packet traverses the service provider's network over an MPLS LSP. It arrives at the egress PE router, which then forwards the traffic to the CE device at the destination customer site.

The key difference in VPLS is that packets can traverse the service provider's network in a point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to PE routers in the VPLS. In contrast, a Layer 2 VPN forwards packets in a point-to-point fashion only. The destination of a packet received from a CE device by a PE router must be known for the Layer 2 VPN to function properly.

VPLS is designed to carry Ethernet traffic across an MPLS-enabled service provider network. In certain ways, VPLS mimics the behavior of an Ethernet network. When a PE router configured with a VPLS routing instance receives a packet from a CE device, it first checks the appropriate routing table for the destination of the VPLS packet. If the router has the destination, it forwards it to the appropriate PE router. If it does not have the destination, it broadcasts the packet to all the other PE routers that are members of the same VPLS routing instance. The PE routers forward the packet to their CE devices. The CE device that is the intended recipient of the packet forwards it to its final destination. The other CE devices discard it.


## Virtual-Router Routing Instances

A virtual-router routing instance, like a VPN routing and forwarding (VRF) routing instance, maintains separate routing and forwarding tables for each instance. However, many configuration steps required for VRF routing instances are not required for virtual-router routing instances. Specifically, you do not need to configure a route distinguisher, a routing table policy (the **vrf-export**, **vrf-import**, and **route-distinguisher** statements), or MPLS between the P routers.

However, you need to configure separate logical interfaces between each of the service provider routers participating in a virtual-router routing instance. You also need to configure separate logical interfaces between the service provider routers and the customer routers participating in each routing instance. Each virtual-router instance requires its own unique set of logical interfaces to all participating routers.

Figure 1 on page 6 shows how this works. The service provider routers G and H are configured for virtual-router routing instances Red and Green. Each service provider router is directly connected to two local customer routers, one in each routing instance. The service provider routers are also connected to each other over the service provider network. These routers need four logical interfaces: a logical interface to each of the locally connected customer routers and a logical interface to carry traffic between the two service provider routers for each virtual-router instance.

*Figure 1: Logical Interface per Router in a Virtual-Router Routing Instance*



Layer 3 VPNs do not have this configuration requirement. If you configure several Layer 3 VPN routing instances on a PE router, all the instances can use the same logical interface to reach another PE router. This is possible because Layer 3 VPNs use MPLS (VPN) labels that differentiate traffic going to and from various routing instances. Without MPLS and VPN labels, as in a virtual-router routing instance, you need separate logical interfaces to separate traffic from different instances.

One method of providing this logical interface between the service provider routers is by configuring tunnels between them. You can configure IP Security (IPsec), generic routing encapsulation (GRE), or IP-IP tunnels between the service provider routers, terminating the tunnels at the virtual-router instance.

## VPNs and Logical Systems

You can partition a single physical router into multiple logical systems that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the physical router, logical systems offer an effective way to maximize the use of a single routing platform.

Logical systems perform a subset of the actions of a physical router and have their own unique routing tables, interfaces, policies, and routing instances. A set of logical systems within a single router can handle the functions previously performed by several small routers.

Logical systems support Layer 2 VPNs, Layer 3 VPNs, VPLS, and Layer 2 circuits.. For more information about logical systems, see the *Logical Systems Feature Guide*.

Starting in Junos OS release 17.4R1, Ethernet VPN (EVPN) support has also been extended to logical systems running on MX devices. The same EVPN options and performance are available, and can be configured under the **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols evpn]** hierarchy.

Release History Table

Release	Description
17.4	Starting in Junos OS release 17.4R1, Ethernet VPN (EVPN) support has also been extended to logical systems running on MX devices. The same EVPN options and performance are available, and can be configured under the <b>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols evpn]</b> hierarchy.

## Layer 2 VPNs

In a Layer 3 network only, you can configure Layer 2 virtual private network (VPN) under a Layer 2 VPN routing instance type **l2vpn**.

In a Layer 2 environment, you can use a **l2vpn** routing instance to transparently carry Layer 2 traffic over an IP/MPLS backbone. Layer 2 traffic is sent to the provider edge (PE) router in Layer 2 format. The PE router encapsulates the frames and transports them over the IP/MPLS backbone to the PE router on the other side of the cloud. The remote PE router removes encapsulation and sends the frames to the receiving site in Layer 2 format.


### Related Documentation

- *MX Series Router Architecture*
- *Layer 2 and Layer 3 Features on MX Series Routers*
- *Junos OS VPNs Library for Routing Devices*

## Routers in a VPN

Figure 2 on page 8 illustrates how VPN functionality is provided by the provider edge (PE) routers; the provider and customer edge (CE) routers have no special configuration requirements for VPNs.

**Figure 2: Routers in a VPN**



## CHAPTER 2

# Assigning Routing Instances to VPNs

- [Configuring Routing Instances on PE Routers in VPNs on page 9](#)
- [Configuring Virtual-Router Routing Instances in VPNs on page 14](#)
- [Configuring Path MTU Checks for VPN Routing Instances on page 16](#)

## Configuring Routing Instances on PE Routers in VPNs

---

You need to configure a routing instance for each VPN on each of the PE routers participating in the VPN. The configuration procedures outlined in this section are applicable to Layer 2 VPNs, Layer 3 VPNs, and VPLS. The configuration procedures specific to each type of VPN are described in the corresponding sections in the other configuration chapters.

To configure routing instances for VPNs, include the following statements:

```
description text;  
instance-type type;  
interface interface-name;  
route-distinguisher (as-number:number | ip-address:number);  
vrf-import [ policy-names ];  
vrf-export [ policy-names ];  
vrf-target {  
    export community-name;  
    import community-name;  
}
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

To configure VPN routing instances, you perform the steps in the following sections:

- [Configuring the Routing Instance Name for a VPN on page 10](#)
- [Configuring the Description on page 10](#)
- [Configuring the Instance Type on page 10](#)
- [Configuring Interfaces for VPN Routing on page 11](#)

- [Configuring the Route Distinguisher on page 13](#)
- [Configuring Automatic Route Distinguishers on page 14](#)

## Configuring the Routing Instance Name for a VPN

The name of the routing instance for a VPN can be a maximum of 128 characters and can contain letters, numbers, and hyphens. In Junos OS Release 9.0 and later, you can no longer specify **default** as the actual routing-instance name. You also cannot use any special characters (! @ # \$ % ^ & \*, + < > : ;) within the name of a routing instance.



**NOTE:** 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.

Specify the routing-instance name with the **routing-instance** statement:

```
routing-instance routing-instance-name {...}
```

You can include this statement at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

## Configuring the Description

To provide a text description for the routing instance, include the **description** statement. If the text includes one or more spaces, enclose them in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show route instance detail** command and has no effect on the operation of the routing instance.

To configure a text description, include the **description** statement:

```
description text;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

## Configuring the Instance Type

The instance type you configure varies depending on whether you are configuring Layer 2 VPNs, Layer 3 VPNs, VPLS, or virtual routers. Specify the instance type by including the **instance-type** statement:

- To enable Layer 2 VPN routing on a PE router, include the **instance-type** statement and specify the value **l2vpn**:

```
instance-type l2vpn;
```

- To enable VPLS routing on a PE router, include the **instance-type** statement and specify the value **vpls**:

```
instance-type vpls;
```

- Layer 3 VPNs require that each PE router have a VPN routing and forwarding (VRF) table for distributing routes within the VPN. To create the VRF table on the PE router, include the **instance-type** statement and specify the value **vrf**:

```
instance-type vrf;
```



**NOTE:** Routing Engine based sampling is not supported on VRF routing instances.

- To enable the virtual-router routing instance, include the **instance-type** statement and specify the value **virtual-router**:

```
instance-type virtual-router;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

## Configuring Interfaces for VPN Routing

On each PE router, you must configure an interface over which the VPN traffic travels between the PE and CE routers.

The sections that follow describe how to configure interfaces for VPNs:

- [General Configuration for VPN Routing on page 11](#)
- [Configuring Interfaces for Layer 3 VPNs on page 12](#)
- [Configuring Interfaces for Carrier-of-Carriers VPNs on page 12](#)
- [Configuring Unicast RPF on VPN Interfaces on page 13](#)

### General Configuration for VPN Routing

The configuration described in this section applies to all types of VPNs. For Layer 3 VPNs and carrier-of-carriers VPNs, complete the configuration described in this section before proceeding to the interface configuration sections specific to those topics.

To configure interfaces for VPN routing, include the **interface** statement:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Specify both the physical and logical portions of the interface name, in the following format:

*physical.logical*

For example, in **at-1/2/1.2**, **at-1/2/1** is the physical portion of the interface name and **2** is the logical portion. If you do not specify the logical portion of the interface name, the value **0** is set by default.

A logical interface can be associated with only one routing instance. If you enable a routing protocol on all instances by specifying **interfaces all** when configuring the master instance of the protocol at the **[edit protocols]** hierarchy level, and if you configure a specific interface for VPN routing at the **[edit routing-instances routing-instance-name]** hierarchy level or at the **[edit logical-systems logical-system-name routing-instances routing-instance-name]** hierarchy level, the latter interface statement takes precedence and the interface is used exclusively for the VPN.

If you explicitly configure the same interface name at the **[edit protocols]** hierarchy level and at either the **[edit routing-instances routing-instance-name]** or **[edit logical-systems logical-system-name routing-instances routing-instance-name]** hierarchy levels, an attempt to commit the configuration fails.

---

### Configuring Interfaces for Layer 3 VPNs

When you configure the Layer 3 VPN interfaces at the **[edit interfaces]** hierarchy level, you must also configure **family inet** when configuring the logical interface:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family inet;
  }
}
```

---

### Configuring Interfaces for Carrier-of-Carriers VPNs

When you configure carrier-of-carriers VPNs, you need to configure the **family mpls** statement in addition to the **family inet** statement for the interfaces between the PE and CE routers. For carrier-of-carriers VPNs, configure the logical interface as follows:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family inet;
    family mpls;
  }
}
```

If you configure **family mpls** on the logical interface and then configure this interface for a non-carrier-of-carriers routing instance, the **family mpls** statement is automatically removed from the configuration for the logical interface, since it is not needed.



### Configuring Unicast RPF on VPN Interfaces

For VPN interfaces that carry IP version 4 or version 6 (IPv4 or IPv6) traffic, you can reduce the impact of denial-of-service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.

You can configure unicast RPF on a VPN interface by enabling unicast RPF on the interface and including the **interface** statement at the **[edit routing-instances routing-instance-name]** hierarchy level.

You cannot configure unicast RPF on the core-facing interfaces. You can only configure unicast RPF on the CE router-to-PE router interfaces on the PE router. However, for virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.

For information about how to configure unicast RPF on VPN interfaces, see *Configuring Unicast RPF*.

### Configuring the Route Distinguisher

Each routing instance that you configure on a PE router must have a unique route distinguisher associated with it. VPN routing instances need a route distinguisher to help BGP to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. If you configure different VPN routing instances with the same route distinguisher, the commit fails.

For Layer 2 VPNs and VPLS, if you have configured the **l2vpn-use-bgp-rules** statement, you must configure a unique route distinguisher for each PE router participating in a specific routing instance.

For other types of VPNs, we recommend that you use a unique route distinguisher for each PE router participating in the routing instance. Although you can use the same route distinguisher on all PE routers for the same VPN routing instance (except for Layer 2 VPNs and VPLS), if you use a unique route distinguisher, you can determine the CE router from which a route originated within the VPN.

To configure a route distinguisher on a PE router, include the **route-distinguisher** statement:

```
route-distinguisher (as-number:number | ip-address:number);
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The route distinguisher is a 6-byte value that you can specify in one of the following formats:

- **as-number:number**, where **as-number** is an autonomous system (AS) number (a 2-byte value) and **number** is any 4-byte value. The AS number can be in the range 1 through 65,535. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the Internet service provider's (ISP's) own or the customer's own AS number.

- ***ip-address:number***, where ***ip-address*** is an IP address (a 4-byte value) and ***number*** is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a nonprivate address in your assigned prefix range.

## Configuring Automatic Route Distinguishers

If you configure the **route-distinguisher-id** statement at the **[edit routing-options]** hierarchy level, a route distinguisher is automatically assigned to the routing instance. If you also configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**.

To assign a route distinguisher automatically, include the **route-distinguisher-id** statement:

```
route-distinguisher-id ip-address;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

A type 1 route distinguisher is automatically assigned to the routing instance using the format ***ip-address:number***. The IP address is specified by the **route-distinguisher-id** statement and the number is unique for the routing instance.

### Related Documentation

- [Configuring Policies for the VRF Table on PE Routers in VPNs on page 25](#)
- [Configuring BGP Route Target Filtering for VPNs on page 37](#)

---

## Configuring Virtual-Router Routing Instances in VPNs

A virtual-router routing instance, like a VRF routing instance, maintains separate routing and forwarding tables for each instance. However, many of the configuration steps required for VRF routing instances are not required for virtual-router routing instances. Specifically, you do not need to configure a route distinguisher, a routing table policy (the **vrf-export**, **vrf-import**, and **route-distinguisher** statements), or MPLS between the service provider routers.

Configure a virtual-router routing instance by including the following statements:

```
description text;  
instance-type virtual-router;  
interface interface-name;  
protocols { ... }
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

The following sections explain how to configure a virtual-router routing instance:

- [Configuring a Routing Protocol Between the Service Provider Routers on page 15](#)
- [Configuring Logical Interfaces Between Participating Routers on page 15](#)

## Configuring a Routing Protocol Between the Service Provider Routers

The service provider routers need to be able to exchange routing information. You can configure the following protocols for the virtual-router routing instance **protocols** statement configuration at the **[edit routing-instances *routing-instance-name*]** hierarchy level:

- BGP
- IS-IS
- LDP
- OSPF
- Protocol Independent Multicast (PIM)
- RIP

You can also configure static routes.

IBGP route reflection is not supported for virtual-router routing instances.

If you configure LDP under a virtual-router instance, LDP routes are placed by default in the routing instance's inet.0 and inet.3 routing tables (for example, sample.inet.0 and sample.inet.3). To restrict LDP routes to only the routing instance's inet.3 table, include the **no-forwarding** statement:

**no-forwarding;**

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols *ldp*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols *ldp*]**

When you restrict the LDP routes to only the inet.3 routing table, the corresponding IGP route in the inet.0 routing table can be redistributed and advertised into other routing protocols.

For information about routing tables, see *Understanding Junos OS Routing Tables*.

## Configuring Logical Interfaces Between Participating Routers

You must configure an interface to each customer router participating in the routing instance and to each P router participating in the routing instance. Each virtual-router routing instance requires its own separate logical interfaces to all P routers participating in the instance. To configure interfaces for virtual-router instances, include the **interface** statement:

**interface** *interface-name*;

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

Specify both the physical and logical portions of the interface name, in the following format:

*physical.logical*

For example, in **at-1/2/1.2**, **at-1/2/1** is the physical portion of the interface name and **2** is the logical portion. If you do not specify the logical portion of the interface name, **0** is set by default.

You must also configure the interfaces at the **[edit interfaces]** hierarchy level.

One method of providing this logical interface between the provider routers is by configuring tunnels between them. You can configure IP Security (IPsec), generic routing encapsulation (GRE), or IP-IP tunnels between the provider routers, terminating the tunnels at the virtual-router instance.

For information about how to configure tunnels and interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

---

## Configuring Path MTU Checks for VPN Routing Instances

---

By default, the maximum transmission unit (MTU) check for VPN routing instances is disabled on M Series routers (except the M320 router) and enabled for the M320 router and T Series routers. On M Series routers, you can configure path MTU checks on the outgoing interfaces for unicast traffic routed on VRF routing instances and on virtual-router routing instances.

When you enable an MTU check, the routing platform sends an Internet Control Message Protocol (ICMP) message when a packet traversing the routing instance exceeds the MTU size and has the **do-not-fragment** bit set. The ICMP message uses the VRF local address as its source address.

For an MTU check to work in a routing instance, you must both include the **vrf-mtu-check** statement at the **[edit chassis]** hierarchy level and assign at least one interface containing an IP address to the routing instance.

For more information about the path MTU check, see the *Junos OS Administration Library*.

To configure path MTU checks, do the tasks described in the following sections:

- [Enabling Path MTU Checks for a VPN Routing Instance on page 17](#)
- [Assigning an IP Address to the VPN Routing Instance on page 17](#)

## Enabling Path MTU Checks for a VPN Routing Instance

To enable path checks on the outgoing interface for unicast traffic routed on a VRF or virtual-router routing instance, include the **vrf-mtu-check** statement at the **[edit chassis]** hierarchy level:

```
[edit chassis]  
vrf-mtu-check;
```

## Assigning an IP Address to the VPN Routing Instance

To ensure that the path MTU check functions properly, at least one IP address must be associated with each VRF or virtual-router routing instance. If an IP address is not associated with the routing instance, ICMP reply messages cannot be sent.

Typically, the VRF or virtual-router routing instance IP address is drawn from among the IP addresses associated with interfaces configured for that routing instance. If none of the interfaces associated with a VRF or virtual-router routing instance is configured with an IP address, you need to explicitly configure a logical loopback interface with an IP address. This interface must then be associated with the routing instance. See [“Configuring Logical Units on the Loopback Interface for Routing Instances in Layer 3 VPNs” on page 137](#) for details.



## CHAPTER 3

# Distributing Routes in VPNs

- [Enabling Routing Information Exchange for VPNs on page 19](#)
- [Configuring IBGP Sessions Between PE Routers in VPNs on page 19](#)
- [Configuring Aggregate Labels for VPNs on page 21](#)
- [Configuring a Signaling Protocol and LSPs for VPNs on page 22](#)
- [Configuring Policies for the VRF Table on PE Routers in VPNs on page 25](#)
- [Configuring the Route Origin for VPNs on page 32](#)

### Enabling Routing Information Exchange for VPNs

---

For Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, VPLS, EVPNs, and Layer 2 circuits to function properly, the service provider's PE and P routers must be able to exchange routing information. For this to happen, you must configure either an IGP (such as OSPF or IS-IS) or static routes on these routers. You configure the IGP on the master instance of the routing protocol process at the **[edit protocols]** hierarchy level, not within the routing instance used for the VPN—that is, not at the **[edit routing-instances]** hierarchy level.

When you configure the PE router, do not configure any summarization of the PE router's loopback addresses at the area boundary. Each PE router's loopback address should appear as a separate route.

- Related Documentation**
- *Example: Configuring IS-IS*
  - *OSPF Feature Guide*

### Configuring IBGP Sessions Between PE Routers in VPNs

---

You must configure an IBGP session between the PE routers to allow the PE routers to exchange information about routes originating and terminating in the VPN. The PE routers rely on this information to determine which labels to use for traffic destined for remote sites.

Configure an IBGP session for the VPN as follows:

```
[edit protocols]
bgp {
```

```
group group-name {  
  type internal;  
  local-address ip-address;  
  family evpn {  
    signaling;  
  }  
  family (inet-vpn | inet6-vpn) {  
    unicast;  
  }  
  family l2vpn {  
    signaling;  
  }  
  neighbor ip-address;  
}
```

The IP address in the **local-address** statement is the address of the loopback interface on the local PE router. The IBGP session for the VPN runs through the loopback address. (You must also configure the loopback interface at the **[edit interfaces]** hierarchy level.)

The IP address in the **neighbor** statement is the loopback address of the neighboring PE router. If you are using RSVP signaling, this IP address is the same address you specify in the **to** statement at the **[edit mpls label-switched-path *lsp-path-name*]** hierarchy level when you configure the MPLS LSP.

The **family** statement allows you to configure the IBGP session for Layer 2 VPNs, VPLS, EVPNs or for Layer 3 VPNs.

- To configure an IBGP session for Layer 2 VPNs and VPLS, include the **signaling** statement at the **[edit protocols bgp group *group-name* family l2vpn]** hierarchy level:

```
[edit protocols bgp group group-name family l2vpn]  
signaling;
```

- To configure an IBGP session for EVPNs, include the **signaling** statement at the **[edit protocols bgp group *group-name* family evpn]** hierarchy level:

```
[edit protocols bgp group group-name family evpn]  
signaling;
```

- To configure an IPv4 IBGP session for Layer 3 VPNs, configure the **unicast** statement at the **[edit protocols bgp group *group-name* family inet-vpn]** hierarchy level:

```
[edit protocols bgp group group-name family inet-vpn]  
unicast;
```

- To configure an IPv6 IBGP session for Layer 3 VPNs, configure the **unicast** statement at the **[edit protocols bgp group *group-name* family inet6-vpn]** hierarchy level:

```
[edit protocols bgp group group-name family inet6-vpn]  
unicast;
```





**NOTE:** You can configure both `family inet` and `family inet-vpn` or both `family inet6` and `family inet6-vpn` within the same peer group. This allows you to enable support for both IPv4 and IPv4 VPN routes or both IPv6 and IPv6 VPN routes within the same peer group.

**Related  
Documentation**

- [Enabling Routing Information Exchange for VPNs on page 19](#)
- [Configuring a Signaling Protocol and LSPs for VPNs on page 22](#)

## Configuring Aggregate Labels for VPNs

Aggregate labels for VPNs allow a Juniper Networks routing platform to aggregate a set of incoming labels (labels received from a peer router) into a single forwarding label that is selected from the set of incoming labels. The single forwarding label corresponds to a single next hop for that set of labels. Label aggregation reduces the number of VPN labels that the router must examine.

For a set of labels to share an aggregate forwarding label, they must belong to the same forwarding equivalence class (FEC). The labeled packets must have the same destination egress interface.

Including the `community community-name` statement with the `aggregate-label` statement lets you specify prefixes with a common origin community. Set by policy on the peer PE, these prefixes represent an FEC on the peer PE router.



**CAUTION:** If the target community is set by mistake instead of the origin community, forwarding problems at the egress PE can result. All prefixes from the peer PE will appear to be in the same FEC, resulting in a single inner label for all CE routers behind a given PE in the same VPN.

To work with route reflectors in Layer 3 VPN networks, the Juniper Networks M10i router aggregates a set of incoming labels only when the routes:

- Are received from the same peer router
- Have the same site of origin community
- Have the same next hop

The next hop requirement is important because route reflectors forward routes originated from different BGP peers to another BGP peer without changing the next hop of those routes.

To configure aggregate labels for VPNs, include the `aggregate-label` statement:

```
aggregate-label {
  community community-name;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary for this statement.

For information about how to configure a community, see *Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions*.

**Related Documentation**

- *Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions*

## Configuring a Signaling Protocol and LSPs for VPNs

For VPNs to function, you must enable a signaling protocol, either the LDP or RSVP on the provider edge (PE) routers and on the provider (P) routers. You also need to configure label-switched paths (LSPs) between the ingress and egress routers. In a typical VPN configuration, you need to configure LSPs from each PE router to all of the other PE routers participating in the VPN in a full mesh.



**NOTE:** As with any configuration involving MPLS, you cannot configure any of the core-facing interfaces on the PE routers over dense Fast Ethernet PICs.

To enable a signaling protocol, perform the steps in one of the following sections:

- [Using LDP for VPN Signaling on page 22](#)
- [Using RSVP for VPN Signaling on page 24](#)

### Using LDP for VPN Signaling

To use LDP for VPN signaling, perform the following steps on the PE and provider (P) routers:

1. Configure LDP on the interfaces in the core of the service provider's network by including the **ldp** statement at the **[edit protocols]** hierarchy level.

You need to configure LDP only on the interfaces between PE routers or between PE and P routers. You can think of these as the “core-facing” interfaces. You do not need to configure LDP on the interface between the PE and customer edge (CE) routers.

```
[edit]
protocols {
  ldp {
    interface type-fpc/pic/port;
  }
}
```

2. Configure the MPLS address family on the interfaces on which you enabled LDP (the interfaces you configured in Step 1) by including the **family mpls** statement at the **[edit interfaces type-fpc/pic/port unit logical-unit-number]** hierarchy level.

```
[edit]
```

```

interfaces {
  type-fpc/pic/port {
    unit logical-unit-number {
      family mpls;
    }
  }
}

```

### 3. Configure OSPF or IS-IS on each PE and P router.

You configure these protocols at the master instance of the routing protocol, not within the routing instance used for the VPN.

- To configure OSPF, include the **ospf** statement at the **[edit protocols]** hierarchy level. At a minimum, you must configure a backbone area on at least one of the router's interfaces.

```

[edit]
protocols {
  ospf {
    area 0.0.0.0 {
      interface type-fpc/pic/port;
    }
  }
}

```

- To configure IS-IS, include the **isis** statement at the **[edit protocols]** hierarchy level and configure the loopback interface and International Organization for Standardization (ISO) family at the **[edit interfaces]** hierarchy level. At a minimum, you must enable IS-IS on the router, configure a network entity title (NET) on one of the router's interfaces (preferably the loopback interface, lo0), and configure the ISO family on all interfaces on which you want IS-IS to run. When you enable IS-IS, Level 1 and Level 2 are enabled by default. The following is the minimum IS-IS configuration. In the **address** statement, **address** is the NET.

```

[edit]
interfaces {
  lo0 {
    unit logical-unit-number {
      family iso {
        address address;
      }
    }
  }
  type-fpc/pic/port {
    unit logical-unit-number {
      family iso;
    }
  }
}
protocols {
  isis {
    interface all;
  }
}

```

For more information about configuring OSPF and IS-IS, see the *OSPF Feature Guide* and *IS-IS Feature Guide*.

## Using RSVP for VPN Signaling

To use RSVP for VPN signaling, perform the following steps:

1. On each PE router, configure traffic engineering.

To do this, you must configure an interior gateway protocol (IGP) that supports traffic engineering (either IS-IS or OSPF) and enable traffic engineering support for that protocol.

To enable OSPF traffic engineering support, include the **traffic-engineering** statement at the **[edit protocols ospf]** hierarchy level:

```
[edit protocols ospf]
traffic-engineering {
  shortcuts;
}
```

For IS-IS, traffic engineering support is enabled by default.

2. On each PE and P router, enable RSVP on the interfaces that participate in the label-switched path (LSP).

On the PE router, these interfaces are the ingress and egress points to the LSP. On the P router, these interfaces connect the LSP between the PE routers. Do not enable RSVP on the interface between the PE and the CE routers, because this interface is not part of the LSP.

To configure RSVP on the PE and P routers, include the **interface** statement at the **[edit protocols rsvp]** hierarchy level. Include one **interface** statement for each interface on which you are enabling RSVP.

```
[edit protocols]
rsvp {
  interface interface-name;
  interface interface-name;
}
```

3. On each PE router, configure an MPLS LSP to the PE router that is the LSP's egress point.

To do this, include the **interface** and **label-switched-path** statements at the **[edit protocols mpls]** hierarchy level:

```
[edit protocols]
mpls {
  interface interface-name;
  label-switched-path path-name {
    to ip-address;
  }
}
```

In the **to** statement, specify the address of the LSP's egress point, which is an address on the remote PE router.

In the **interface** statement, specify the name of the interface (both the physical and logical portions). Include one **interface** statement for the interface associated with the LSP.

When you configure the logical portion of the same interface at the **[edit interfaces]** hierarchy level, you must also configure the **family inet** and **family mpls** statements:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family inet;
    family mpls;
  }
}
```

4. On all P routers that participate in the LSP, enable MPLS by including the **interface** statement at the **[edit mpls]** hierarchy level.

Include one **interface** statement for each connection to the LSP.

```
[edit]
mpls {
  interface interface-name;
  interface interface-name;
}
```

5. Enable MPLS on the interface between the PE and CE routers by including the **interface** statement at the **[edit mpls]** hierarchy level.

Doing this allows the PE router to assign an MPLS label to traffic entering the LSP or to remove the label from traffic exiting the LSP.

```
[edit]
mpls {
  interface interface-name;
}
```

For information about configuring MPLS, see the *Configuring the Ingress Router for MPLS-Signaled LSPs*.

**Related Documentation**

- [Configuring the Ingress Router for MPLS-Signaled LSPs](#)

## Configuring Policies for the VRF Table on PE Routers in VPNs

On each PE router, you must define policies that define how routes are imported into and exported from the router's VRF table. In these policies, you must define the route target, and you can optionally define the route origin.

To configure policy for the VRF tables, you perform the steps in the following sections:

- [Configuring the Route Target on page 26](#)
- [Configuring the Route Origin on page 26](#)
- [Configuring an Import Policy for the PE Router's VRF Table on page 27](#)

- [Configuring an Export Policy for the PE Router's VRF Table on page 29](#)
- [Applying Both the VRF Export and the BGP Export Policies on page 30](#)
- [Configuring a VRF Target on page 31](#)

## Configuring the Route Target

As part of the policy configuration for the VPN routing table, you must define a route target, which defines which VPN the route is a part of. When you configure different types of VPN services (Layer 2 VPNs, Layer 3 VPNs, EVPNs, or VPLS) on the same PE router, be sure to assign unique route target values to avoid the possibility of adding route and signaling information to the wrong VPN routing table.

To configure the route target, include the **target** option in the **community** statement:

```
community name members target:community-id;
```

You can include this statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

*name* is the name of the community.

*community-id* is the identifier of the community. Specify it in one of the following formats:

- ***as-number:number***, where *as-number* is an AS number (a 2-byte value) and *number* is a 4-byte community value. The AS number can be in the range 1 through 65,535. We recommend that you use an IANA-assigned, nonprivate AS number, preferably the ISP's own or the customer's own AS number. The community value can be a number in the range 0 through 4,294,967,295 ( $2^{32} - 1$ ).
- ***ip-address:number***, where *ip-address* is an IPv4 address (a 4-byte value) and *number* is a 2-byte community value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a nonprivate address in your assigned prefix range. The community value can be a number in the range 1 through 65,535.

## Configuring the Route Origin

In the import and export policies for the PE router's VRF table, you can optionally assign the route origin (also known as the site of origin) for a PE router's VRF routes using a VRF export policy applied to multiprotocol external BGP (MP-EBGP) VPN IPv4 route updates sent to other PE routers.

Matching on the assigned route origin attribute in a receiving PE's VRF import policy helps ensure that VPN-IPv4 routes learned through MP-EBGP updates from one PE are not reimported to the same VPN site from a different PE connected to the same site.

To configure a route origin, complete the following steps:

1. Include the **community** statement with the **origin** option:

```
community name members origin:community-id;
```

You can include this statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

**name** is the name of the community.

**community-id** is the identifier of the community. Specify it in one of the following formats:

- **as-number:number**, where **as-number** is an AS number (a 2-byte value) and **number** is a 4-byte community value. The AS number can be in the range 1 through 65,535. We recommend that you use an IANA-assigned, nonprivate AS number, preferably the ISP's own or the customer's own AS number. The community value can be a number in the range 0 through 4,294,967,295 ( $2^{32} - 1$ ).
  - **ip-address:number**, where **ip-address** is an IPv4 address (a 4-byte value) and **number** is a 2-byte community value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a nonprivate address in your assigned prefix range. The community value can be a number in the range 1 through 65,535.
2. Include the community in the import policy for the PE router's VRF table by configuring the **community** statement with the **community-id** identifier defined in Step 1 at the **[edit policy-options policy-statement import-policy-name term import-term-name from]** hierarchy level. See [“Configuring an Import Policy for the PE Router's VRF Table” on page 27](#).  
  
If the policy's **from** clause does not specify a community condition, the **vrf-import** statement in which the policy is applied cannot be committed. The Junos OS commit operation does not pass the validation check.
  3. Include the community in the export policy for the PE router's VRF table by configuring the **community** statement with the **community-id** identifier defined in Step 1 at the **[edit policy-options policy-statement export-policy-name term export-term-name then]** hierarchy level. See [“Configuring an Export Policy for the PE Router's VRF Table” on page 29](#).

See [“Configuring the Route Origin for VPNs” on page 32](#) for a configuration example.

## Configuring an Import Policy for the PE Router's VRF Table

Each VPN can have a policy that defines how routes are imported into the PE router's VRF table. An import policy is applied to routes received from other PE routers in the VPN. A policy must evaluate all routes received over the IBGP session with the peer PE router. If the routes match the conditions, the route is installed in the PE router's **routing-instance-name.inet.0** VRF table. An import policy must contain a second term that rejects all other routes.

Unless an import policy contains only a **then reject** statement, it must include a reference to a community. Otherwise, when you try to commit the configuration, the commit fails. You can configure multiple import policies.

An import policy determines what to import to a specified VRF table based on the VPN routes learned from the remote PE routers through IBGP. The IBGP session is configured at the **[edit protocols bgp]** hierarchy level. If you also configure an import policy at the **[edit protocols bgp]** hierarchy level, the import policies at the **[edit policy-options]** hierarchy level and the **[edit protocols bgp]** hierarchy level are combined through a logical AND operation. This allows you to filter traffic as a group.

To configure an import policy for the PE router's VRF table, follow these steps:

1. To define an import policy, include the **policy-statement** statement. For all PE routers, an import policy must always include the **policy-statement** statement, at a minimum:

```
policy-statement import-policy-name {  
  term import-term-name {  
    from {  
      protocol bgp;  
      community community-id;  
    }  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}
```

You can include the **policy-statement** statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

The ***import-policy-name*** policy evaluates all routes received over the IBGP session with the other PE router. If the routes match the conditions in the **from** statement, the route is installed in the PE router's *routing-instance-name*.inet.0 VRF table. The second term in the policy rejects all other routes.

For more information about creating policies, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

2. You can optionally use a regular expression to define a set of communities to be used for the VRF import policy.

For example you could configure the following using the **community** statement at the **[edit policy-options policy-statement *policy-statement-name*]** hierarchy level:

```
[edit policy-options vrf-import-policy-sample]  
community high-priority members *:50
```

Note that you cannot configure a regular expression as a part of a route target extended community. For more information about how to configure regular expressions for



communities, see *Understanding How to Define BGP Communities and Extended Communities*.

3. To configure an import policy, include the **vrf-import** statement:

```
vrf-import import-policy-name;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

## Configuring an Export Policy for the PE Router's VRF Table

Each VPN can have a policy that defines how routes are exported from the PE router's VRF table. An export policy is applied to routes sent to other PE routers in the VPN. An export policy must evaluate all routes received over the routing protocol session with the CE router. (This session can use the BGP, OSPF, or Routing Information Protocol [RIP] routing protocols, or static routes.) If the routes match the conditions, the specified community target (which is the route target) is added to them and they are exported to the remote PE routers. An export policy must contain a second term that rejects all other routes.

Export policies defined within the VPN routing instance are the only export policies that apply to the VRF table. Any export policy that you define on the IBGP session between the PE routers has no effect on the VRF table. You can configure multiple export policies.

To configure an export policy for the PE router's VRF table, follow these steps:

1. For all PE routers, an export policy must distribute VPN routes to and from the connected CE routers in accordance with the type of routing protocol that you configure between the CE and PE routers within the routing instance.

To define an export policy, include the **policy-statement** statement. An export policy must always include the **policy-statement** statement, at a minimum:

```
policy-statement export-policy-name {
  term export-term-name {
    from protocol (bgp | ospf | rip | static);
    then {
      community add community-id;
      accept;
    }
  }
  term term-name {
    then reject;
  }
}
```



**NOTE:** Configuring the `community add` statement is a requirement for Layer 2 VPN VRF export policies. If you change the `community add` statement to the `community set` statement, the router at the egress of the Layer 2 VPN link might drop the connection.



**NOTE:** When configuring draft-rosen multicast VPNs operating in source-specific mode and using the `vrf-export` statement to specify the export policy, the policy must have a term that accepts routes from the `vrf-name.mdt.0` routing table. This term ensures proper PE autodiscovery using the `inet-mdt` address family.

When configuring draft-rosen multicast VPNs operating in source-specific mode and using the `vrf-target` statement, the VRF export policy is automatically generated and automatically accepts routes from the `vrf-name.mdt.0` routing table.

You can include the `policy-statement` statement at the following hierarchy levels:

- [edit `policy-options`]
- [edit `logical-systems logical-system-name policy-options`]

The `export-policy-name` policy evaluates all routes received over the routing protocol session with the CE router. (This session can use the BGP, OSPF, or RIP routing protocols, or static routes.) If the routes match the conditions in the `from` statement, the community target specified in the `then community add` statement is added to them and they are exported to the remote PE routers. The second term in the policy rejects all other routes.

For more information about creating policies, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

2. To apply the policy, include the `vrf-export` statement:

```
vrf-export export-policy-name;
```

You can include this statement at the following hierarchy levels:

- [edit `routing-instances routing-instance-name`]
- [edit `logical-systems logical-system-name routing-instances routing-instance-name`]

## Applying Both the VRF Export and the BGP Export Policies

When you apply a VRF export policy as described in “[Configuring an Export Policy for the PE Router’s VRF Table](#)” on page 29, routes from VPN routing instances are advertised to other PE routers based on this policy, whereas the BGP export policy is ignored.

If you include the **vpn-apply-export** statement in the BGP configuration, both the VRF export and BGP group or neighbor export policies are applied (VRF first, then BGP) before routes are advertised in the VPN routing tables to other PE routers.

When you include the **vpn-apply-export** statement, be aware of the following:

- Routes imported into the `l3vpn.bgp.0` routing table retain the attributes of the original routes (for example, an OSPF route remains an OSPF route even when it is stored in the `l3vpn.bgp.0` routing table). You should be aware of this when you configure an export policy for connections between an IBGP PE router and a PE router, a route reflector and a PE router, or AS boundary router (ASBR) peer routers.
- By default, all routes in the `l3vpn.bgp.0` routing table are exported to the IBGP peers. If the last statement of the export policy is deny all and if the export policy does not specifically match on routes in the `l3vpn.bgp.0` routing table, no routes are exported.

To apply both the VRF export and BGP export policies to VPN routes, include the **vpn-apply-export** statement:

```
vpn-apply-export;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring a VRF Target

Including the **vrf-target** statement in the configuration for a VRF target community causes default VRF import and export policies to be generated that accept and tag routes with the specified target community. You can still create more complex policies by explicitly configuring VRF import and export policies. These policies override the default policies generated when you configure the **vrf-target** statement.

If you do not configure the **import** and **export** options of the **vrf-target** statement, the specified community string is applied in both directions. The **import** and **export** keywords give you more flexibility, allowing you to specify a different community for each direction.

The syntax for the VRF target community is not a name. You must specify it in the format **target:x.y**. A community name cannot be specified because this would also require you to configure the community members for that community using the **policy-options** statement. If you define the **policy-options** statements, then you can just configure VRF import and export policies as usual. The purpose of the **vrf-target** statement is to simplify the configuration by allowing you to configure most statements at the **[edit routing-instances]** hierarchy level.

To configure a VRF target, include the **vrf-target** statement:

```
vrf-target community;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

An example of how you might configure the **vrf-target** statement follows:

```
[edit routing-instances sample]
vrf-target target:69:102;
```

To configure the **vrf-target** statement with the **export** and **import** options, include the following statements:

```
vrf-target {
  export community-name;
  import community-name;
}
```

You can include this statement at the following hierarchy levels:


- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

## Configuring the Route Origin for VPNs

You can use route origin to prevent routes learned from one customer edge (CE) router marked with origin community from being advertised back to it from another CE router in the same AS.

In the example, the route origin is used to prevent routes learned from CE Router A that are marked with origin community from being advertised back to CE Router E by AS 200. The example topology is shown in [Figure 3 on page 32](#).

**Figure 3: Network Topology of Site of Origin Example**



In this topology, CE Router A and CE Router E are in the same AS (AS200). They use EBGP to exchange routes with their respective provider edge (PE) routers, PE Router B and PE Router D. The two CE routers have a back connection.

The following sections describe how to configure the route origin for a group of VPNs:

- [Configuring the Site of Origin Community on CE Router A on page 33](#)
- [Configuring the Community on CE Router A on page 33](#)
- [Applying the Policy Statement on CE Router A on page 33](#)
- [Configuring the Policy on PE Router D on page 34](#)
- [Configuring the Community on PE Router D on page 34](#)
- [Applying the Policy on PE Router D on page 35](#)

## Configuring the Site of Origin Community on CE Router A

The following section describes how to configure CE Router A to advertise routes with a site of origin community to PE Router B for this example.



**NOTE:** In this example, direct routes are configured to be advertised, but any route can be configured.

Configure a policy to advertise routes with **my-soo** community on CE Router A as follows:

```
[edit]
policy-options {
  policy-statement export-to-my-isp {
    term a {
      from {
        protocol direct;
      }
      then {
        community add my-soo;
        accept;
      }
    }
  }
}
```

## Configuring the Community on CE Router A

Configure the **my-soo** community on CE Router A as follows:

```
[edit]
policy-options {
  community my-soo {
    members origin:100:1;
  }
}
```

## Applying the Policy Statement on CE Router A

Apply the export-to-my-isp policy statement as an export policy to the EBGp peering on the CE Router A as follows:

```
[edit]
protocols {
  bgp {
    group my_isp {
      export export-to-my-isp;
    }
  }
}
```

When you issue the **show route receive-protocol bgp detail** command, you should see the following routes originated from PE Router B with **my-soo** community:

```
user@host> show route receive-protocol bgp 10.12.99.2 detail
inet.0: 16 destinations, 16 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpn_blue.inet.0: 8 destinations, 10 routes (8 active, 0 holddown, 0 hidden)
* 10.12.33.0/30 (2 entries, 1 announced)
    Nexthop: 10.12.99.2
    AS path: 100 I
    Communities: origin:100:1
10.12.99.0/30 (2 entries, 1 announced)
    Nexthop: 10.12.99.2
    AS path: 100 I
    Communities: origin:100:1
* 10.255.71.177/32 (1 entry, 1 announced)
    Nexthop: 10.12.99.2
    AS path: 100 I
    Communities: origin:100:1
* 192.168.64.0/21 (1 entry, 1 announced)
    Nexthop: 10.12.99.2
    AS path: 100 I
    Communities: origin:100:1
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
```

## Configuring the Policy on PE Router D

Configure a policy on PE Router D that prevents routes with **my-soo** community tagged by CE Router A from being advertised to CE Router E as follows:

```
[edit]
policy-options {
  policy-statement soo-ce1-policy {
    term a {
      from {
        community my-soo;
      }
      then {
        reject;
      }
    }
  }
}
```

## Configuring the Community on PE Router D

Configure the community on PE Router D as follows:

```
[edit]
policy-options {
  community my-soo {
    members origin:100:1;
  }
}
```

```

    }
}

```

## Applying the Policy on PE Router D

To prevent routes learned from CE Router A from being advertised to CE Router E (the two routers can communicate these routes directly), apply the **soo-ce1-policy** policy statement as an export policy to the PE Router D and CE Router E EBGp session **vpn\_blue**.

View the EBGp session on PE Router D using the **show routing-instances** command.

```

user@host# show routing-instances
vpn_blue {
  instance-type vrf;
  interface fe-2/0/0.0;
  vrf-target target:100:200;
  protocols {
    bgp {
      group ce2 {
        advertise-peer-as;
        peer-as 100;
        neighbor 10.12.99.6;
      }
    }
  }
}

```

Apply the **soo-ce1-policy** policy statement as an export policy to the PE Router D and CE Router E EBGp session **vpn\_blue** as follows:

```

[edit routing-instances]
vpn_blue {
  protocols {
    bgp {
      group ce2 {
        export soo-ce1-policy;
      }
    }
  }
}

```





## CHAPTER 4

# Distributing VPN Routes with Target Filtering

- [Configuring BGP Route Target Filtering for VPNs on page 37](#)
- [Example: BGP Route Target Filtering for VPNs on page 39](#)
- [Example: Configuring BGP Route Target Filtering for VPNs on page 41](#)
- [Configuring Static Route Target Filtering for VPNs on page 50](#)
- [Understanding Proxy BGP Route Target Filtering for VPNs on page 50](#)
- [Example: Configuring Proxy BGP Route Target Filtering for VPNs on page 51](#)
- [Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs on page 67](#)
- [Reducing Network Resource Use with Static Route Target Filtering for VPNs on page 84](#)

## Configuring BGP Route Target Filtering for VPNs

---

BGP route target filtering allows you to distribute VPN routes to only the routers that need them. In VPN networks without BGP route target filtering configured, BGP distributes all VPN routes to all VPN peer routers.

For more information about BGP route target filtering, see RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*.

The following sections provide an overview of BGP route target filtering and how to configure it for VPNs:

- [BGP Route Target Filtering Overview on page 37](#)
- [Configuring BGP Route Target Filtering for VPNs on page 38](#)

## BGP Route Target Filtering Overview

PE routers, unless they are configured as route reflectors or are running an EBGp session, discard any VPN routes that do not include a route target extended community as specified in the local VRF import policies. This is the default behavior of the Junos OS.

However, unless it is explicitly configured not to store VPN routes, any router configured either as a route reflector or border router for a VPN address family must store all of the

VPN routes that exist in the service provider's network. Also, though PE routers can automatically discard routes that do not include a route target extended community, route updates continue to be generated and received.

By reducing the number of routers receiving VPN routes and route updates, BGP route target filtering helps to limit the amount of overhead associated with running a VPN. BGP route target filtering is most effective at reducing VPN-related administrative traffic in networks where there are many route reflectors or AS border routers that do not participate in the VPNs directly (not acting as PE routers for the CE devices).

BGP route target filtering uses standard UPDATE messages to distribute route target extended communities between routers. The use of UPDATE messages allows BGP to use its standard loop detection mechanisms, path selection, policy support, and database exchange implementation.

## Configuring BGP Route Target Filtering for VPNs

BGP route target filtering is enabled through the exchange of the **route-target** address family, stored in the `bgp.rtarget.0` routing table. Based on the **route-target** address family, the route target NLRI (address family indicator [AFI]=1, subsequent AFI [SAFI]=132) is negotiated with its peers.

On a system that has locally configured VRF instances, BGP automatically generates local routes corresponding to targets referenced in the **vrf-import** policies.

To configure BGP route target filtering, include the **family route-target** statement:

```
family route-target {  
    advertise-default;  
    external-paths number;  
    prefix-limit number;  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The **advertise-default**, **external-paths**, and **prefix-limit** statements affect the BGP route target filtering configuration as follows:

- The **advertise-default** statement causes the router to advertise the default route target route (0:0:0/0) and suppress all routes that are more specific. This can be used by a route reflector on BGP groups consisting of neighbors that act as PE routers only. PE routers often need to advertise all routes to the route reflector.

Suppressing all route target advertisements other than the default route reduces the amount of information exchanged between the route reflector and the PE routers. The Junos OS further helps to reduce route target advertisement overhead by not maintaining dependency information unless a nondefault route is received.

- The **external-paths** statement (which has a default value of 1) causes the router to advertise the VPN routes that reference a given route target. The number you specify determines the number of external peer routers (currently advertising that route target) that receive the VPN routes.

- The **prefix-limit** statement limits the number of prefixes that can be received from a peer router.

The **route-target**, **advertise-default**, and **external-path** statements affect the RIB-OUT state and must be consistent between peer routers that share the same BGP group. The **prefix-limit** statement affects the receive side only and can have different settings between different peer routers in a BGP group.

#### Related Documentation

- [Understanding Proxy BGP Route Target Filtering for VPNs on page 50](#)
- [Example: Configuring Proxy BGP Route Target Filtering for VPNs on page 51](#)
- [Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs on page 67](#)
- [Example: Configuring BGP Route Target Filtering for VPNs on page 41](#)
- [Configuring the Route Origin for VPNs on page 32](#)
- [Configuring Static Route Target Filtering for VPNs on page 50](#)

## Example: BGP Route Target Filtering for VPNs

BGP route target filtering is enabled by configuring the **family route-target** statement at the appropriate BGP hierarchy level. This statement enables the exchange of a new **route-target** address family, which is stored in the `bgp.target.0` routing table.

The following configuration illustrates how you could configure BGP route target filtering for a BGP group titled **to\_vpn04**:

```
[edit]
protocols {
  bgp {
    group to_vpn04 {
      type internal;
      local-address 10.255.14.182;
      peer-as 200;
      neighbor 10.255.14.174 {
        family inet-vpn {
          unicast;
        }
        family route-target;
      }
    }
  }
}
```

The following configuration illustrates how you could configure a couple of local VPN routing and forwarding (VRF) routing instances to take advantage of the functionality provided by BGP route target filtering. Based on this configuration, BGP would automatically generate local routes corresponding to the route targets referenced in the VRF import policies (note the targets defined by the **vrf-target** statements).

```
[edit]
```

```

routing-instances {
  vpn1 {
    instance-type vrf;
    interface t1-0/1/2.0;
    vrf-target target:200:101;
    protocols {
      ospf {
        export bgp-routes;
        area 0.0.0.0 {
          interface t1-0/1/2.0;
        }
      }
    }
  }
  vpn2 {
    instance-type vrf;
    interface t1-0/1/2.1;
    vrf-target target:200:102;
    protocols {
      ospf {
        export bgp-routes;
        area 0.0.0.0 {
          interface t1-0/1/2.1;
        }
      }
    }
  }
}

```

Issue the **show route table bgp.rtarget.0** show command to verify the BGP route target filtering configuration:

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 6 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
200:200:101/96
    *[RTarget/5] 00:10:00
        Local
200:200:102/96
    *[RTarget/5] 00:10:00
        Local
200:200:103/96
    *[BGP/170] 00:09:48, localpref 100, from 10.255.14.174
        AS path: I
        > t3-0/0/0.0
200:200:104/96
    *[BGP/170] 00:09:48, localpref 100, from 10.255.14.174
        AS path: I
        > t3-0/0/0.0

```

The **show** command display format for route target prefixes is:

*AS number:route target extended community/length*

The first number represents the autonomous system (AS) of the router that sent this advertisement. The remainder of the display follows the Junos **show** command convention for extended communities.

The output from the **show route table bgp-rtarget.0** command displays the locally generated and remotely generated routes.

The first two entries correspond to the route targets configured for the two local VRF routing instances (**vpn1** and **vpn2**):

- **200:200:101/96**—Community **200:101** in the **vpn1** routing instance
- **200:200:102/96**—Community **200:102** in the **vpn2** routing instance

The last two entries are prefixes received from a BGP peer:


- **200:200:103/96**—Tells the local router that routes tagged with this community (**200:103**) should be advertised to peer **10.255.14.174** through **t3-0/0/0.0**
- **200:200:104/96**—Tells the local router that routes tagged with this community (**200:104**) should be advertised to peer **10.255.14.174** through **t3-0/0/0.0**

## Example: Configuring BGP Route Target Filtering for VPNs

BGP route target filtering reduces the number of routers that receive VPN routes and route updates, helping to limit the amount of overhead associated with running a VPN. BGP route target filtering is most effective at reducing VPN-related administrative traffic in networks where there are many route reflectors or AS border routers that do not participate in the VPNs directly (do not act as PE routers for the CE devices).

Figure 4 on page 41 illustrates the topology for a network configured with BGP route target filtering for a group of VPNs.

*Figure 4: BGP Route Target Filtering Enabled for a Group of VPNs*



The following sections describe how to configure BGP route target filtering for a group of VPNs:

- [Configure BGP Route Target Filtering on Router PE1 on page 42](#)
- [Configure BGP Route Target Filtering on Router PE2 on page 43](#)
- [Configure BGP Route Target Filtering on the Route Reflector on page 46](#)
- [Configure BGP Route Target Filtering on Router PE3 on page 47](#)

## Configure BGP Route Target Filtering on Router PE1

This section describes how to enable BGP route target filtering on Router PE1 for this example.

Configure the routing options on router PE1 as follows:

```
[edit]
routing-options {
  route-distinguisher-id 10.255.14.182;
  autonomous-system 198;
}
```

Configure the BGP protocol on Router PE1 as follows:

```
[edit]
protocols {
  bgp {
    group to_VPN_D {
      type internal;
      local-address 10.255.14.182;
      peer-as 198;
      neighbor 10.255.14.174 {
        family inet-vpn {
          unicast;
        }
        family route-target;
      }
    }
  }
}
```

Configure the **vpn1** routing instance as follows:

```
[edit]
routing-instances {
  vpn1 {
    instance-type vrf;
    interface t1-0/1/2.0;
    vrf-target target:198:101;
    protocols {
      ospf {
        export bgp-routes;
        area 0.0.0.0 {
          interface t1-0/1/2.0;
        }
      }
    }
  }
}
```

Configure the **vpn2** routing instance on Router PE1 as follows:

```
[edit]
routing-instances {
  vpn2 {
```

```

instance-type vrf;
interface t1-0/1/2.1;
vrf-target target:198:102;
protocols {
  ospf {
    export bgp-routes;
    area 0.0.0.0 {
      interface t1-0/1/2.1;
    }
  }
}
}
}

```

Once you have implemented this configuration, you should see the following when you issue a **show route table bgp.rtarget.0** command:

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 6 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

198.198:101/96
    *[RTarget/5] 00:27:42
        Local
        [BGP/170] 00:27:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/0.0
198.198:102/96
    *[RTarget/5] 00:27:42
        Local
        [BGP/170] 00:27:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/0.0
198.198:103/96
    *[BGP/170] 00:27:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/0.0
198.198:104/96
    *[BGP/170] 00:27:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/0.0

```

## Configure BGP Route Target Filtering on Router PE2

This section describes how to enable BGP route target filtering on Router PE2 for this example.

Configure the routing options on Router PE2 as follows:

```

[edit]
routing-options {
  route-distinguisher-id 10.255.14.176;
  autonomous-system 198;
}

```

Configure the BGP protocol on Router PE2 as follows:

```
[edit]
protocols {
  bgp {
    group to_vpn04 {
      type internal;
      local-address 10.255.14.176;
      peer-as 198;
      neighbor 10.255.14.174 {
        family inet-vpn {
          unicast;
        }
        family route-target;
      }
    }
  }
}
```

Configure the **vpn1** routing instance on Router PE2 as follows:

```
[edit]
routing-instances {
  vpn1 {
    instance-type vrf;
    interface t3-0/0/0.0;
    vrf-target target:198:101;
    protocols {
      bgp {
        group vpn1 {
          type external;
          peer-as 101;
          as-override;
          neighbor 10.49.11.2;
        }
      }
    }
  }
}
```

Configure the **vpn2** routing instance on Router PE2 as follows:

```
[edit]
routing-instances {
  vpn2 {
    instance-type vrf;
    interface t3-0/0/0.1;
    vrf-target target:198:102;
    protocols {
      bgp {
        group vpn2 {
          type external;
          peer-as 102;
          as-override;
          neighbor 10.49.21.2;
        }
      }
    }
  }
}
```



```

    }
  }
}

```

Configure the **vpn3** routing instance on Router PE2 as follows:

```

[edit]
routing-instances {
  vpn3 {
    instance-type vrf;
    interface t3-0/0/0.2;
    vrf-import vpn3-import;
    vrf-export vpn3-export;
    protocols {
      bgp {
        group vpn3 {
          type external;
          peer-as 103;
          as-override;
          neighbor 10.49.31.2;
        }
      }
    }
  }
}

```

Once you have configured router PE2 in this manner, you should see the following when you issue the **show route table bgp.rtarget.0** command:

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 7 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

198.198:101/96
    *[RTarget/5] 00:28:15
        Local
        [BGP/170] 00:28:03, localpref 100, from
10.255.14.174
    AS path: I
    > via t1-0/1/0.0
198.198:102/96
    *[RTarget/5] 00:28:15
        Local
        [BGP/170] 00:28:03, localpref 100, from
10.255.14.174
    AS path: I
    > via t1-0/1/0.0
198.198:103/96
    *[RTarget/5] 00:28:15
        Local
        [BGP/170] 00:28:03, localpref 100, from
10.255.14.174
    AS path: I
    > via t1-0/1/0.0
198.198:104/96
    *[BGP/170] 00:28:03, localpref 100, from
10.255.14.174
    AS path: I
    > via t1-0/1/0.0

```

## Configure BGP Route Target Filtering on the Route Reflector

This section illustrates how to enable BGP route target filtering on the route reflector for this example.

Configure the routing options on the route reflector as follows:

```
[edit]
routing-options {
  route-distinguisher-id 10.255.14.174;
  autonomous-system 198;
}
```

Configure the BGP protocol on the route reflector as follows:

```
[edit]
protocols {
  bgp {
    group rr-group {
      type internal;
      local-address 10.255.14.174;
      cluster 10.255.14.174;
      peer-as 198;
      neighbor 10.255.14.182 {
        description to_PE1_vpn12;
        family inet-vpn {
          unicast;
        }
        family route-target;
      }
      neighbor 10.255.14.176 {
        description to_PE2_vpn06;
        family inet-vpn {
          unicast;
        }
        family route-target;
      }
      neighbor 10.255.14.178 {
        description to_PE3_vpn08;
        family inet-vpn {
          unicast;
        }
        family route-target;
      }
    }
  }
}
```

Once you have configured the route reflector in this manner, you should see the following when you issue the **show route table bgp.rtarget.0** command:

```
user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 8 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

198.198:101/96
```

```

10.255.14.176      *[BGP/170] 00:29:03, localpref 100, from
                   AS path: I
                   > via t1-0/2/0.0
10.255.14.178      [BGP/170] 00:29:03, localpref 100, from
                   AS path: I
                   > via t3-0/1/1.0
10.255.14.182      [BGP/170] 00:29:03, localpref 100, from
                   AS path: I
                   > via t3-0/1/3.0
198.198:102/96     *[BGP/170] 00:29:03, localpref 100, from
10.255.14.176      AS path: I
                   > via t1-0/2/0.0
10.255.14.182      [BGP/170] 00:29:03, localpref 100, from
                   AS path: I
                   > via t3-0/1/3.0
198.198:103/96     *[BGP/170] 00:29:03, localpref 100, from
10.255.14.176      AS path: I
                   > via t1-0/2/0.0
10.255.14.178      [BGP/170] 00:29:03, localpref 100, from
                   AS path: I
                   > via t3-0/1/1.0
198.198:104/96     *[BGP/170] 00:29:03, localpref 100, from
10.255.14.178      AS path: I
                   > via t3-0/1/1.0

```

### Configure BGP Route Target Filtering on Router PE3

The following section describes how to enable BGP route target filtering on Router PE3 for this example.

Configure the routing options on Router PE3 as follows:

```

[edit]
routing-options {
  route-distinguisher-id 10.255.14.178;
  autonomous-system 198;
}

```

Configure the BGP protocol on Router PE3 as follows:

```

[edit]
protocols {
  bgp {
    group to_vpn04 {
      type internal;
      local-address 10.255.14.178;
      peer-as 198;
      neighbor 10.255.14.174 {

```

```
        family inet-vpn {
            unicast;
        }
        family route-target;
    }
}
}
```

Configure the **vpn1** routing instance on Router PE3 as follows:

```
[edit]
routing-instances {
  vpn1 {
    instance-type vrf;
    interface t3-0/0/0.0;
    vrf-target target:198:101;
    protocols {
      rip {
        group vpn1 {
          export bgp-routes;
          neighbor t3-0/0/0.0;
        }
      }
    }
  }
}
```

Configure the **vpn3** routing instance on Router PE3 as follows:

```
[edit]
routing-instances {
  vpn3 {
    instance-type vrf;
    interface t3-0/0/0.1;
    vrf-target target:198:103;
    protocols {
      rip {
        group vpn3 {
          export bgp-routes;
          neighbor t3-0/0/0.1;
        }
      }
    }
  }
}
```

Configure the **vpn4** routing instance on Router PE3 as follows:

```
[edit]
routing-instances {
  vpn4 {
    instance-type vrf;
    interface t3-0/0/0.2;
    vrf-target target:198:104;
    protocols {
      rip {
```

```

    group vpn4 {
      export bgp-routes;
      neighbor t3-0/0/0.2;
    }
  }
}

```

Once you have configured Router PE3 in this manner, you should see the following when you issue the **show route table bgp.rtarget.0** command:

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 7 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

198.198:101/96
    *[RTarget/5] 00:29:42
        Local
        [BGP/170] 00:29:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/1.0
198.198:102/96
    *[BGP/170] 00:29:29, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/1.0
198.198:103/96
    *[RTarget/5] 00:29:42
        Local
        [BGP/170] 00:29:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/1.0
198.198:104/96
    *[RTarget/5] 00:29:42
        Local
        [BGP/170] 00:29:30, localpref 100, from
10.255.14.174
    AS path: I
    > via t3-0/0/1.0

```

## Configuring Static Route Target Filtering for VPNs

---

The BGP VPN route target extended community (RFC 4360, *BGP Extended Communities Attribute*) is used to determine VPN membership. Static route target filtering helps to prevent resources from being consumed in portions of the network where the VPN routes are not needed due to the lack of member PE routers (RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*). Routers can originate routes into the RT-Constrain protocol to indicate their interest in receiving VPN routes containing route targets that match the RT-Constrain NLRI.

To configure static route target filtering for VPNs:

- Configure the **route-target-filter** statement at the **[edit routing-options rib bgp.rtarget.0 static]** hierarchy level.

The following example illustrates how you could configure the **route-target-filter** statement:

```
[edit routing-options rib bgp.rtarget.0 static]
route-target-filter destination {
  group bgp-group;
  local;
  neighbor bgp-peer;
}
```

- You can display route target filtering information using the **show bgp group rtf detail** command.

### Related Documentation

- [Reducing Network Resource Use with Static Route Target Filtering for VPNs on page 84](#)
- [Configuring BGP Route Target Filtering for VPNs on page 37](#)

## Understanding Proxy BGP Route Target Filtering for VPNs

---

BGP route target filtering (also known as route target constrain, or RTC) allows you to distribute VPN routes to only the devices that need them. In VPN networks without BGP route target filtering configured, BGP distributes all VPN routes to all VPN peer devices, which can strain network resources. The route target filtering feature was introduced to reduce the number of devices receiving VPN routes and VPN routing updates, thereby limiting the amount of overhead associated with running a VPN. The Junos OS implementation for BGP route target filtering is based on RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*.

What if you have a network environment where route target filtering is not widely deployed, or what if some devices do not support route target filtering? For example, you might have a BGP speaker with route target filtering enabled that is peered with a BGP speaker that does not support or have route target filtering configured. In this case, the BGP speaker with route target filtering configured must advertise default route target

membership (RT membership) on behalf of its peer. The route target filtering resource savings are unrealized because the device supporting the filtering must now send all VPN routes to the device that does not support the filter. Proxy BGP route target filtering (or Proxy RTC) permits the generation of RT membership for devices that do not support route target filtering. This eases the deployment of route target filtering in networks where it is incompletely deployed or not fully supported.

Proxy BGP route target filtering allows you to distribute proxy RT membership advertisements created from the received BGP VPN routes to other devices in the network that need them. These are known as proxy advertisements because the device creates the RT membership on behalf of its peers without the route target filtering functionality. Proxy BGP route target filtering uses BGP route target extended communities that are exported to a specific BGP speaker to generate the route targets. Generated proxy RTC routes are stored in the `bgp.target.0` routing table.

You can also configure a policy to control which VPN routes are used to generate the proxy RTC routes. This can help control which RT membership is generated by the proxying device. In addition, you can configure a policy to reduce the memory overhead associated with proxy RTC. Proxy RTC only uses additional memory on a per-VPN route basis when it is permitted by a policy to be used for generating RT membership.

#### Related Documentation

- [Configuring BGP Route Target Filtering for VPNs on page 37](#)
- [Example: Configuring Proxy BGP Route Target Filtering for VPNs on page 51](#)
- [Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs on page 67](#)

## Example: Configuring Proxy BGP Route Target Filtering for VPNs

This example shows how to configure proxy BGP route target filtering (also known as proxy route target constrain, or proxy RTC).

- [Requirements on page 51](#)
- [Overview on page 52](#)
- [Configuration on page 53](#)
- [Verification on page 66](#)

### Requirements

This example uses the following hardware and software components:

- Four Juniper Networks devices that can be a combination of M Series, MX Series, or T Series routers.
- Junos OS Release 12.2 or later on one or more devices configured for proxy BGP route filtering. In this example, you explicitly configure proxy BGP route filtering on the route reflectors.

Before configuring proxy BGP route target filtering, make sure that you are familiar with and understand the following concepts:

- [Layer 2 VPNs](#)
- [Layer 3 VPNs on page 123](#)
- [Route distinguishers on page 217](#)
- [Route targets on page 25](#)
- [BGP route target filtering on page 37](#)
- [BGP extended communities](#)

## Overview

Route target filtering decreases the number of devices in a network that receive VPN routes that are not needed. Proxy BGP route target filtering allows networks to take advantage of route target filtering in locations where the feature is not currently supported. By configuring this feature, you can realize many of the same network resource savings that are available to you if your network fully supported BGP route target filtering.

To configure proxy BGP route target filtering, you include the **family route-target proxy-generate** statement on the devices that will distribute proxy route target membership (RT membership) advertisements for the devices that do not support BGP route target filtering. The proxy BGP route target filtering routes are then stored in the `bgp.rtarget.0` routing table.

Proxy BGP route target filtering is intended to create RT membership advertisements for devices that do not support the BGP route target filtering feature. If the **proxy-generate** statement is present, but the route target family is negotiated with the BGP peer, the proxy-generate functionality is disabled. This allows simplified configuration of BGP peer groups where a portion of the peers in the group support route target filtering but others do not. In such an example case, the **family route-target proxy-generate** statement might be part of the BGP peer group configuration.



.....

**NOTE:** When deploying proxy BGP route target filtering in your network, the **advertise-default** statement for BGP route target filtering causes the device to advertise the default route target route (0:0:0/0) and suppress all routes that are more specific. If you have proxy BGP route target filtering configured on one device and one or more peers have the **advertise-default** statement configured as part of their BGP route target filtering configuration, the **advertise-default** configuration is ignored.

.....


---

## Topology Diagram

Figure 5 on page 53 shows the topology used in this example.



Figure 5: Proxy BGP Route Target Filtering Topology



In this example, BGP route target filtering is configured on the route reflectors (Device RR1 and Device RR2) and the provider edge (PE) Device PE2, but the other PE, Device PE1, does not support the BGP route target filtering functionality. Device PE2 has four VPNs configured (vpn1, vpn2, vpn3, and vpn4). Device PE1 has two VPNs configured (vpn1 and vpn2), so this device is only interested in receiving route updates for vpn1 and vpn2. Currently, this is impossible because both route reflectors (Device RR1 and Device RR2) learn and share information about all of the incoming VPN routes (vpn1 through vpn4) with Device PE1. In the sample topology, all devices participate in autonomous system (AS) 203, OSPF is the configured interior gateway protocol (IGP), and LDP is the signaling protocol used by the VPNs. In this example, we use static routes in the VPN routing and forwarding (VRF) instances to generate VPN routes. This is done in place of using a PE to customer edge (CE) protocol such as OSPF or BGP.

To minimize the number of VPN route updates being processed by Device PE1, you include the **family route-target proxy-generate** statement to configure proxy BGP route target filtering on each route reflector. Each route reflector has a peering session with Device PE1 and supports route target filtering to the core. However, Device PE1 does not support route target filtering, so the network resource savings are unrealized by Device PE1 since it receives all of the VPN updates. By configuring proxy BGP route target filtering on the peering sessions facing Device PE1, you limit the number of VPN updates processed by Device PE1, and the route reflectors generate the proxy BGP route target routes for Device PE1 throughout the network.

## Configuration

- [Configuring Device PE1 on page 55](#)
- [Configuring Device RR1 on page 58](#)
- [Configuring Device RR2 on page 60](#)
- [Configuring Device PE2 on page 63](#)

### 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 PE1
set interfaces ge-1/0/0 unit 0 description PE1-to-RR1
set interfaces ge-1/0/0 unit 0 family inet address 10.49.0.1/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description PE1-to-RR2
```

```

set interfaces ge-1/0/1 unit 0 family inet address 10.49.10.1/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.255.163.58
set protocols bgp group internal neighbor 10.255.165.220 family inet-vpn unicast
set protocols bgp group internal neighbor 10.255.165.28 family inet-vpn unicast
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.163.58
set routing-options autonomous-system 203
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 vrf-target target:203:100
set routing-instances vpn1 routing-options static route 203.0.113.1/24 discard
set routing-instances vpn2 instance-type vrf
set routing-instances vpn2 vrf-target target:203:101
set routing-instances vpn2 routing-options static route 203.0.113.2/24 discard

```

**Device RR1**

```

set interfaces ge-1/0/0 unit 0 description RR1-to-PE1
set interfaces ge-1/0/0 unit 0 family inet address 10.49.0.2/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description RR1-to-PE2
set interfaces ge-1/0/1 unit 0 family inet address 10.50.0.2/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 198.51.100.1
set protocols bgp group internal cluster 198.51.100.1
set protocols bgp group internal neighbor 10.255.163.58 description vpn1-to-pe1 family
inet-vpn unicast
set protocols bgp group internal neighbor 10.255.163.58 family route-target proxy-generate
set protocols bgp group internal neighbor 10.255.168.42 description vpn1-to-pe2 family
inet-vpn unicast
set protocols bgp group internal neighbor 10.255.168.42 family route-target
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.165.220
set routing-options autonomous-system 203

```

**Device RR2**

```

set interfaces ge-1/0/0 unit 0 description RR2-to-PE1
set interfaces ge-1/0/0 unit 0 family inet address 10.49.10.2/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description RR2-to-PE2
set interfaces ge-1/0/1 unit 0 family inet address 10.50.10.2/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.255.165.28
set protocols bgp group internal cluster 198.51.100.1

```

```

set protocols bgp group internal neighbor 10.255.163.58 description vpn2-to-pe1 family
  inet-vpn unicast
set protocols bgp group internal neighbor 10.255.163.58 family route-target proxy-generate
set protocols bgp group internal neighbor 10.255.168.42 description vpn2-to-pe2 family
  inet-vpn unicast
set protocols bgp group internal neighbor 10.255.168.42 family route-target
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.165.28
set routing-options autonomous-system 203

```

**Device PE2**

```

set interfaces ge-1/0/0 unit 0 description PE2-to-RR1
set interfaces ge-1/0/0 unit 0 family inet address 10.50.0.1/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description PE2-to-RR2
set interfaces ge-1/0/1 unit 0 family inet address 10.50.10.1/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.255.168.42
set protocols bgp group internal family inet-vpn unicast
set protocols bgp group internal family route-target
set protocols bgp group internal neighbor 10.255.165.220
set protocols bgp group internal neighbor 10.255.165.28
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.168.42
set routing-options autonomous-system 203
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 vrf-target target:203:100
set routing-instances vpn1 routing-options static route 203.0.113.1/24 discard
set routing-instances vpn2 instance-type vrf
set routing-instances vpn2 vrf-target target:203:101
set routing-instances vpn2 routing-options static route 203.0.113.2/24 discard
set routing-instances vpn3 instance-type vrf
set routing-instances vpn3 vrf-target target:203:103
set routing-instances vpn3 routing-options static route 203.0.113.3/24 discard
set routing-instances vpn4 instance-type vrf
set routing-instances vpn4 vrf-target target:203:104
set routing-instances vpn4 routing-options static route 203.0.113.4/24 discard

```

**Configuring Device PE1****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*.

To configure Device PE1:

1. Configure the interfaces.

[edit interfaces]

```
user@PE1# set ge-1/0/0 unit 0 description PE1-to-RR1
user@PE1# set ge-1/0/0 unit 0 family inet address 10.49.0.1/30
user@PE1# set ge-1/0/0 unit 0 family mpls
```

```
user@PE1# set ge-1/0/1 unit 0 description PE1-to-RR2
user@PE1# set ge-1/0/1 unit 0 family inet address 10.49.10.1/30
user@PE1# set ge-1/0/1 unit 0 family mpls
```

2. Configure the route distinguisher and the AS number.

```
[edit routing-options]
user@PE1# set route-distinguisher-id 10.255.163.58
user@PE1# set autonomous-system 203
```

3. Configure LDP as the signaling protocol used by the VPN.

```
[edit protocols ldp]
user@PE1# set interface ge-1/0/0
user@PE1# set interface ge-1/0/1
```

4. Configure BGP.

```
[edit protocols bgp group internal]
user@PE1# set type internal
user@PE1# set local-address 10.255.163.58
user@PE1# set neighbor 10.255.165.220 family inet-vpn unicast
user@PE1# set neighbor 10.255.165.28 family inet-vpn unicast
```

5. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface ge-1/0/0
user@PE1# set interface ge-1/0/1
user@PE1# set interface lo0.0 passive
```

6. Configure the VPN routing instances.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set vrf-target target:203:100
user@PE1# set routing-options static route 203.0.113.1/24 discard

[edit routing-instances vpn2]
user@PE1# set instance-type vrf
user@PE1# set vrf-target target:203:101
user@PE1# set routing-options static route 203.0.113.2/24 discard
```

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

```
[edit]
user@PE1# commit
```

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

```

user@PE1# show interfaces
ge-1/0/0 {
  unit 0 {
    description PE1-to-RR1;
    family inet {
      address 10.49.0.1/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description PE1-to-RR2;
    family inet {
      address 10.49.10.1/30;
    }
    family mpls;
  }
}

user@PE1# show protocols
bgp {
  group internal {
    type internal;
    local-address 10.255.163.58;
    neighbor 10.255.165.220 {
      family inet-vpn {
        unicast;
      }
    }
    neighbor 10.255.165.28 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

```

```

user@PE1# show routing-options
route-distinguisher-id 10.255.14.182;
autonomous-system 203;

user@PE1# show routing-instances
vpn1 {
  instance-type vrf;
  vrf-target target:203:100;
  routing-options {
    static {
      route 203.0.113.1/24 discard;
    }
  }
}
vpn2 {
  instance-type vrf;
  vrf-target target:203:101;
  routing-options {
    static {
      route 203.0.113.2/24 discard;
    }
  }
}

```

### Configuring Device RR1

**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*.

To configure Device RR1:

1. Configure the interfaces.

```

[edit interfaces]
user@RR1# set ge-1/0/0 unit 0 description RR1-to-PE1
user@RR1# set ge-1/0/0 unit 0 family inet address 10.49.0.2/30
user@RR1# set ge-1/0/0 unit 0 family mpls

```

```

user@RR1# set ge-1/0/1 unit 0 description RR1-to-PE2
user@RR1# set ge-1/0/1 unit 0 family inet address 10.50.0.2/30
user@RR1# set ge-1/0/1 unit 0 family mpls

```

2. Configure the route distinguisher and the AS number.

```

[edit routing-options]
user@RR1# set route-distinguisher-id 10.255.165.220
user@RR1# set autonomous-system 203

```

3. Configure LDP as the signaling protocol used by the VPN.

```

[edit protocols ldp]
user@RR1# set interface ge-1/0/0
user@RR1# set interface ge-1/0/1

```

## 4. Configure BGP.

```
[edit protocols bgp group internal]
user@RR1# set type internal
user@RR1# set local-address 10.255.165.220
user@RR1# set cluster 198.51.100.1
user@RR1# set neighbor 10.255.163.58 description vpn1-to-pe1 family inet-vpn
unicast
user@RR1# set neighbor 10.255.168.42 description vpn1-to-pe2 family inet-vpn
unicast
```

## 5. Configure BGP route target filtering on the peering session with Device PE2.

```
[edit protocols bgp group internal]
user@RR1# set neighbor 10.255.168.42 family route-target
```

## 6. Configure proxy BGP route target filtering on the peering session with Device PE1.

```
[edit protocols bgp group internal]
user@RR1# set neighbor 10.255.163.58 family route-target proxy-generate
```

## 7. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@RR1# set interface ge-1/0/0
user@RR1# set interface ge-1/0/1
user@RR1# set interface lo0.0 passive
```

## 8. If you are done configuring the device, commit the configuration.

```
[edit]
user@RR1# commit
```

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

```
user@RR1# show interfaces
ge-1/0/0 {
  unit 0 {
    description RR1-to-PE1;
    family inet {
      address 10.49.0.2/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description RR1-to-PE2;
    family inet {
      address 10.50.0.2/30;
    }
  }
}
```

```
        family mpls;
    }
}

user@RR1# show protocols
bgp {
  group internal {
    type internal;
    local-address 198.51.100.1;
    cluster 198.51.100.1;
    neighbor 10.255.163.58 {
      description vpn1-to-pe1;
      family inet-vpn {
        unicast;
      }
      family route-target {
        proxy-generate;
      }
    }
    neighbor 10.255.168.42 {
      description vpn1-to-pe2;
      family inet-vpn {
        unicast;
      }
      family route-target;
    }
  }
}

ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}

ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@RR1# show routing-options
route-distinguisher-id 10.255.165.220;
autonomous-system 203;
```

---

### Configuring Device RR2

#### 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*.

To configure Device RR2:

1. Configure the interfaces.



```
[edit interfaces]
user@RR2# set ge-1/0/0 unit 0 description RR2-to-PE1
user@RR2# set ge-1/0/0 unit 0 family inet address 10.49.10.2/30
user@RR2# set ge-1/0/0 unit 0 family mpls
```

```
user@RR2# set ge-1/0/1 unit 0 description RR2-to-PE2
user@RR2# set ge-1/0/1 unit 0 family inet address 10.50.10.2/30
user@RR2# set ge-1/0/1 unit 0 family mpls
```

2. Configure the route distinguisher and the AS number.

```
[edit routing-options]
user@RR2# set route-distinguisher-id 10.255.165.28
user@RR2# set autonomous-system 203
```

3. Configure LDP as the signaling protocol used by the VPN.

```
[edit protocols ldp]
user@RR2# set interface ge-1/0/0
user@RR2# set interface ge-1/0/1
```

4. Configure BGP.

```
[edit protocols bgp group internal]
user@RR2# set type internal
user@RR2# set local-address 10.255.165.28
user@RR2# set cluster 198.51.100.1
user@RR2# set neighbor 10.255.163.58 description vpn2-to-pe1 family inet-vpn
unicast
user@RR2# set neighbor 10.255.168.42 description vpn2-to-pe2 family inet-vpn
unicast
```

5. Configure BGP route target filtering on the peering session with Device PE2.

```
[edit protocols bgp group internal]
user@RR2# set neighbor 10.255.168.42 family route-target
```

6. Configure proxy BGP route target filtering on the peering session with Device PE1.

```
[edit protocols bgp group internal]
user@RR2# set neighbor 10.255.163.58 family route-target proxy-generate
```

7. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@RR2# set interface ge-1/0/0
user@RR2# set interface ge-1/0/1
user@RR2# set interface lo0.0 passive
```

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

```
[edit]
user@RR2# commit
```

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

```
user@RR2# show interfaces
ge-1/0/0 {
  unit 0 {
    description RR2-to-PE1;
    family inet {
      address 10.49.10.2/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description RR2-to-PE2;
    family inet {
      address 10.50.10.2/30;
    }
    family mpls;
  }
}

user@RR2# show protocols
bgp {
  group internal {
    local-address 10.255.165.28;
    cluster 198.51.100.1;
    neighbor 10.255.163.58 {
      description vpn2-to-pe1;
      family inet-vpn {
        unicast;
      }
      family route-target {
        proxy-generate;
      }
    }
    neighbor 10.255.168.42 {
      description vpn2-to-pe2;
      family inet-vpn {
        unicast;
      }
      family route-target;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
```

```

ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@RR2# show routing-options
route-distinguisher-id 10.255.165.28;
autonomous-system 203;

```

### Configuring Device PE2

**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*.

To configure Device PE2:

1. Configure the interfaces.

```

[edit interfaces]
user@PE2# set ge-1/0/0 unit 0 description PE2-to-RR1
user@PE2# set ge-1/0/0 unit 0 family inet address 10.50.0.1/30
user@PE2# set ge-1/0/0 unit 0 family mpls

user@PE2# set ge-1/0/1 unit 0 description PE2-to-RR2
user@PE2# set ge-1/0/1 unit 0 family inet address 10.50.10.1/30
user@PE2# set ge-1/0/1 unit 0 family mpls

```

2. Configure the route distinguisher and the AS number.

```

[edit routing-options]
user@PE2# set route-distinguisher-id 10.255.168.42
user@PE2# set autonomous-system 203

```

3. Configure LDP as the signaling protocol used by the VPN.

```

[edit protocols ldp]
user@PE2# set interface ge-1/0/0
user@PE2# set interface ge-1/0/1

```

4. Configure BGP.

```

[edit protocols bgp group internal]
user@PE2# set type internal
user@PE2# set local-address 10.255.168.42
user@PE2# set family inet-vpn unicast
user@PE2# set family route-target
user@PE2# set neighbor 10.255.165.220
user@PE2# set neighbor 10.255.165.28

```

5. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@PE2# set interface ge-1/0/0

```

```

user@PE2# set interface ge-1/0/1
user@PE2# set interface lo0.0 passive

```

6. Configure the VPN routing instances.

```

[edit routing-instances vpn1]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:100
user@PE2# set routing-options static route 203.0.113.1/24 discard

[edit routing-instances vpn2]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:101
user@PE2# set routing-options static route 203.0.113.2/24 discard

[edit routing-instances vpn3]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:103
user@PE2# set routing-options static route 203.0.113.3/24 discard

[edit routing-instances vpn4]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:104
user@PE2# set routing-options static route 203.0.113.4/24 discard

```

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

```

[edit]
user@PE2# commit

```

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

```

user@PE2# show interfaces
ge-1/0/0 {
  unit 0 {
    description PE2-to-RR1;
    family inet {
      address 10.50.0.1/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description PE2-to-RR2;
    family inet {
      address 10.50.10.1/30;
    }
    family mpls;
  }
}

```

```
user@PE2# show protocols
bgp {
  group internal {
    type internal;
    local-address 10.255.168.42;
    family inet-vpn {
      unicast;
    }
    family route-target;
    neighbor 10.255.165.220;
    neighbor 10.255.165.28;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@PE2# show routing-options
route-distinguisher-id 10.255.168.42;
autonomous-system 203;

user@PE2# show routing-instances
vpn1 {
  instance-type vrf;
  vrf-target target:203:100;
  routing-options {
    static {
      route 203.0.113.1/24 discard;
    }
  }
}
vpn2 {
  instance-type vrf;
  vrf-target target:203:101;
  routing-options {
    static {
      route 203.0.113.2/24 discard;
    }
  }
}
vpn3 {
  instance-type vrf;
  vrf-target target:203:103;
  routing-options {
    static {
      route 203.0.113.3/24 discard;
    }
  }
}
```

```

    }
  }
  vpn4 {
    instance-type vrf;
    vrf-target target:203:104;
    routing-options {
      static {
        route 203.0.113.4/24 discard;
      }
    }
  }
}

```

## Verification

Confirm that the configuration is working properly.

### Verifying the Proxy BGP Route Target Routes

**Purpose** Verify that the proxy BGP route target routes are displayed in the `bgp.rtarget.0` table on Device RR1.

**Action** From operational mode, enter the `show route table bgp.rtarget.0` command to display the proxy BGP route targets.

```

user@RR1# show route table bgp.rtarget.0
4 destinations, 6 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

203:203:100/96
    * [RTarget/5] 00:01:22
      Type Proxy
      for 10.255.163.58
      Local
      [BGP/170] 00:04:55, localpref 100, from 10.255.168.42
      AS path: I, validation-state: unverified
      > to 10.50.0.1 via ge-1/0/1

203:203:101/96
    * [RTarget/5] 00:01:22
      Type Proxy
      for 10.255.163.58
      Local
      [BGP/170] 00:04:55, localpref 100, from 10.255.168.42
      AS path: I, validation-state: unverified
      > to 10.50.0.1 via ge-1/0/1

203:203:103/96
    * [BGP/170] 00:04:55, localpref 100, from 10.255.168.42
      AS path: I, validation-state: unverified
      > to 10.50.0.1 via ge-1/0/1

203:203:104/96
    * [BGP/170] 00:04:55, localpref 100, from 10.255.168.42
      AS path: I, validation-state: unverified
      > to 10.50.0.1 via ge-1/0/1

```

**Meaning** Device RR1 is generating the proxy BGP route target routes on behalf of its peer Device PE1. The proxy BGP route target routes are identified with the protocol and preference [RTarget/5] and the route target type of **Proxy**.

- Related Documentation**
- [Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs on page 67](#)
  - [Configuring BGP Route Target Filtering for VPNs on page 37](#)
  - [Understanding Proxy BGP Route Target Filtering for VPNs on page 50](#)

## Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs

This example shows how to configure an export routing policy for BGP route target filtering (also known as route target constrain, or RTC).

- [Requirements on page 67](#)
- [Overview on page 67](#)
- [Configuration on page 69](#)
- [Verification on page 83](#)

### Requirements

This example uses the following hardware and software components:

- Four Juniper Networks devices that support BGP route target filtering.
- Junos OS Release 12.2 or later on one or more devices configured for proxy BGP route filtering. In this example, you explicitly configure proxy BGP route filtering on the route reflectors.

Before configuring an export policy for BGP route target filtering, make sure that you are familiar with and understand the following concepts:

- *Layer 2 VPNs*
- [Layer 3 VPNs on page 123](#)
- [Route distinguishers on page 217](#)
- [Route targets on page 25](#)
- [BGP route target filtering on page 37](#)
- *BGP extended communities*

### Overview

BGP route target filtering allows you to reduce network resource consumption by distributing route target membership (RT membership) advertisements throughout the network. BGP uses the RT membership information to send VPN routes only to the devices that need them in the network. Similar to other types of BGP reachability, you can apply

a routing policy to route target filtering routes to influence the network. When route target filtering is configured, restricting the flow of route target filtering routes also restricts the VPN routes that might be attracted by this RT membership. Configuring this policy involves:

- Creating a filter that defines the list of route target prefixes.
- Creating a policy to select a subset of the route target filters to use for BGP route target filtering.

To define the list of route target prefixes:

- You configure the **rtf-prefix-list** statement at the **[edit policy-options]** hierarchy level to specify the name of the route target prefix list and one or more route target prefixes to use. This configuration allows you to specify the incoming route target filtering routes that the device will use and then distribute them throughout the network.

To configure the routing policy and apply the route target prefix list to that policy, you can specify the following policy options:

- **family route-target**—(Optional) The route-target family match condition specifies matching BGP route target filtering routes. You define this criteria in the **from** statement. This example shows how to create an export policy using the **family route-target** match condition.
- **protocol route-target**—(Optional) The route-target protocol match condition defines the criteria that an incoming route must match. You define this criteria in the **from** statement. This statement is primarily useful for restricting the policy to locally generated route target filtering routes.



**NOTE:** When you use the `show route table bgp.rtarget.0` command to view proxy BGP route target filtering routes, you will see the BGP protocol for received routes and the route target protocol routes for local route target filtering routes.

---

- **rtf-prefix-list name**—The **rtf-prefix-list** statement applies the list of route target prefixes that you already configured to the policy. You define this criteria in the **from** statement.


---

### Topology Diagram

Figure 6 on page 69 shows the topology used in this example.



Figure 6: BGP Route Target Filtering Export Policy Topology



In this example, BGP route target filtering is configured on the route reflectors (Device RR1 and Device RR2) and provider edge (PE) Device PE2. The other PE, Device PE1, does not support BGP route target filtering. Proxy BGP route target filtering is also configured on the peering sessions between the route reflectors and Device PE1 to minimize the number of VPN route updates processed by Device PE1. Device PE2 has four VPNs configured (vpn1, vpn2, vpn3, and vpn4), and Device PE1 has two VPNs configured (vpn1 and vpn2). In the sample topology, all devices participate in autonomous system (AS) 203, OSPF is the configured interior gateway protocol (IGP), and LDP is the signaling protocol used by the VPNs. In this example, we use static routes in the VPN routing and forwarding (VRF) instances to generate VPN routes. This is done in place of using a PE to customer edge (CE) protocol such as OSPF or BGP.

In this example, you further control the routes being advertised from Device PE2 to Device PE1 by configuring an export policy on Device PE2 to prevent vpn3 routes from being advertised to Device RR1. You create a policy that specifies the **family route-target** match condition, defines the list of route target prefixes, and applies the list of route target prefixes by defining the **rtf-prefix-list** criteria.

## Configuration

- [Configuring Device PE1 on page 71](#)
- [Configuring Device RR1 on page 74](#)
- [Configuring Device RR2 on page 77](#)
- [Configuring Device PE2 on page 79](#)

**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 PE1
set interfaces ge-1/0/0 unit 0 description PE1-to-RR1
set interfaces ge-1/0/0 unit 0 family inet address 10.49.0.1/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description PE1-to-RR2
set interfaces ge-1/0/1 unit 0 family inet address 10.49.10.1/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.255.163.58
```

```
set protocols bgp group internal neighbor 10.255.165.220 family inet-vpn unicast
set protocols bgp group internal neighbor 10.255.165.28 family inet-vpn unicast
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.163.58
set routing-options autonomous-system 203
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 vrf-target target:203:100
set routing-instances vpn1 routing-options static route 203.0.113.1/24 discard
set routing-instances vpn2 instance-type vrf
set routing-instances vpn2 vrf-target target:203:101
set routing-instances vpn2 routing-options static route 203.0.113.2/24 discard
```

**Device RR1**

```
set interfaces ge-1/0/0 unit 0 description RR1-to-PE1
set interfaces ge-1/0/0 unit 0 family inet address 10.49.0.2/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description RR1-to-PE2
set interfaces ge-1/0/1 unit 0 family inet address 10.50.0.2/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 198.51.100.0
set protocols bgp group internal cluster 198.51.100.1
set protocols bgp group internal neighbor 10.255.163.58 description vpn1-to-pe1 family
  inet-vpn unicast
set protocols bgp group internal neighbor 10.255.163.58 family route-target proxy-generate
set protocols bgp group internal neighbor 10.255.168.42 description vpn1-to-pe2 family
  inet-vpn unicast
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.165.220
set routing-options autonomous-system 203
```

**Device RR2**

```
set interfaces ge-1/0/0 unit 0 description RR2-to-PE1
set interfaces ge-1/0/0 unit 0 family inet address 10.49.10.2/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description RR2-to-PE2
set interfaces ge-1/0/1 unit 0 family inet address 10.50.10.2/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.255.165.28
set protocols bgp group internal cluster 198.51.100.1
set protocols bgp group internal neighbor 10.255.163.58 description vpn2-to-pe1 family
  inet-vpn unicast
set protocols bgp group internal neighbor 10.255.163.58 family route-target proxy-generate
set protocols bgp group internal neighbor 10.255.168.42 description vpn2-to-pe2 family
  inet-vpn unicast
set protocols bgp group internal neighbor 10.255.163.58 family route-target
set protocols ospf area 0.0.0.0 interface ge-1/0/0
```

```

set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-options route-distinguisher-id 10.255.165.28
set routing-options autonomous-system 203

```

**Device PE2**

```

set interfaces ge-1/0/0 unit 0 description PE2-to-RR1
set interfaces ge-1/0/0 unit 0 family inet address 10.50.0.1/30
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 description PE2-to-RR2
set interfaces ge-1/0/1 unit 0 family inet address 10.50.10.2/30
set interfaces ge-1/0/1 unit 0 family mpls
set protocols ldp interface ge-1/0/0
set protocols ldp interface ge-1/0/1
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.255.168.42
set protocols bgp group internal family inet-vpn unicast
set protocols bgp group internal family route-target
set protocols bgp group internal neighbor 10.255.165.220 export filter-rtc
set protocols bgp group internal neighbor 10.255.165.28
set protocols ospf area 0.0.0.0 interface ge-1/0/0
set protocols ospf area 0.0.0.0 interface ge-1/0/1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options rtf-prefix-list exclude-103 203:203:103/96
set policy-options policy-statement filter-rtc from family route-target
set policy-options policy-statement filter-rtc from rtf-prefix-list exclude-103
set policy-options policy-statement filter-rtc then reject
set routing-options route-distinguisher-id 10.255.168.42
set routing-options autonomous-system 203
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 vrf-target target:203:100
set routing-instances vpn1 routing-options static route 203.0.113.1/24 discard
set routing-instances vpn2 instance-type vrf
set routing-instances vpn2 vrf-target target:203:101
set routing-instances vpn2 routing-options static route 203.0.113.2/24 discard
set routing-instances vpn3 instance-type vrf
set routing-instances vpn3 vrf-target target:203:103
set routing-instances vpn3 routing-options static route 203.0.113.3/24 discard
set routing-instances vpn4 instance-type vrf
set routing-instances vpn4 vrf-target target:203:104
set routing-instances vpn4 routing-options static route 203.0.113.4/24 discard

```

### Configuring Device PE1

---

**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*.

To configure Device PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set ge-1/0/0 unit 0 description PE1-to-RR1
user@PE1# set ge-1/0/0 unit 0 family inet address 10.49.0.1/30
user@PE1# set ge-1/0/0 unit 0 family mpls

```

```
user@PE1#set ge-1/0/1 unit 0 description PE1-to-RR2
user@PE1#set ge-1/0/1 unit 0 family inet address 10.49.10.1/30
user@PE1# set ge-1/0/1 unit 0 family mpls
```

2. Configure the route distinguisher and the AS number.

```
[edit routing-options]
user@PE1# set route-distinguisher-id 10.255.163.58
user@PE1# set autonomous-system 203
```

3. Configure LDP as the signaling protocol used by the VPN.

```
[edit protocols ldp]
user@PE1# set interface ge-1/0/0
user@PE1# set interface ge-1/0/1
```

4. Configure BGP.

```
[edit protocols bgp group internal]
user@PE1# set type internal
user@PE1# set local-address 10.255.163.58
user@PE1# set neighbor 10.255.165.220 family inet-vpn unicast
user@PE1# set neighbor 10.255.165.28 family inet-vpn unicast
```

5. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface ge-1/0/0
user@PE1# set interface ge-1/0/1
user@PE1# set interface lo0.0 passive
```

6. Configure the VPN routing instances.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set vrf-target target:203:100
user@PE1# set routing-options static route 203.0.113.1/24 discard

[edit routing-instances vpn2]
user@PE1# set instance-type vrf
user@PE1# set vrf-target target:203:101
user@PE1# set routing-options static route 203.0.113.2/24 discard
```

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

```
[edit]
user@PE1# commit
```

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

```
user@PE1# show interfaces
ge-1/0/0 {
  unit 0 {
    description PE1-to-RR1;
    family inet {
      address 10.49.0.1/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description PE1-to-RR2;
    family inet {
      address 10.49.10.1/30;
    }
    family mpls;
  }
}

user@PE1# show protocols
bgp {
  group internal {
    type internal;
    local-address 10.255.163.58;
    neighbor 10.255.165.220 {
      family inet-vpn {
        unicast;
      }
    }
    neighbor 10.255.165.28 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@PE1# show routing-options
route-distinguisher-id 10.255.14.182;
autonomous-system 203;

user@PE1# show routing-instances
vpn1 {
```

```

instance-type vrf;
vrf-target target:203:100;
routing-options {
  static {
    route 203.0.113.1/24 discard;
  }
}
}
vpn2 {
  instance-type vrf;
  vrf-target target:203:101;
  routing-options {
    static {
      route 203.0.113.2/24 discard;
    }
  }
}
}

```

### Configuring Device RR1

**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*.

To configure Device RR1:

1. Configure the interfaces.

```

[edit interfaces]
user@RR1# set ge-1/0/0 unit 0 description RR1-to-PE1
user@RR1# set ge-1/0/0 unit 0 family inet address 10.49.0.2/30
user@RR1# set ge-1/0/0 unit 0 family mpls
user@RR1# set ge-1/0/1 unit 0 description RR1-to-PE2
user@RR1# set ge-1/0/1 unit 0 family inet address 10.50.0.2/30
user@RR1# set ge-1/0/1 unit 0 family mpls

```

2. Configure the route distinguisher and the AS number.

```

[edit routing-options]
user@RR1# set route-distinguisher-id 10.255.165.220
user@RR1# set autonomous-system 203

```

3. Configure LDP as the signaling protocol used by the VPN.

```

[edit protocols ldp]
user@RR1# set interface ge-1/0/0
user@RR1# set interface ge-1/0/1

```

4. Configure BGP.

```

[edit protocols bgp group internal]
user@RR1# set type internal
user@RR1# set local-address 10.255.165.220
user@RR1# set cluster 198.51.100.1

```

```

user@RR1# set neighbor 10.255.163.58 description vpn1-to-pe1 family inet-vpn
unicast
user@RR1# set neighbor 10.255.168.42 description vpn1-to-pe2 family inet-vpn
unicast

```

5. Configure BGP route target filtering on the peering session with Device PE2.

```

[edit protocols bgp group internal]
user@RR1# set neighbor 10.255.168.42 family route-target

```

6. Configure proxy BGP route target filtering on the peering session with Device PE1.

```

[edit protocols bgp group internal]
user@RR1# set neighbor 10.255.163.58 family route-target proxy-generate

```

7. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@RR1# set interface ge-1/0/0
user@RR1# set interface ge-1/0/1
user@RR1# set interface lo0.0 passive

```

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

```

[edit]
user@RR1# commit

```

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

```

user@RR1# show interfaces
ge-1/0/0 {
  unit 0 {
    description RR1-to-PE1;
    family inet {
      address 10.49.0.2/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description RR1-to-PE2;
    family inet {
      address 10.50.0.2/30;
    }
    family mpls;
  }
}

user@RR1# show protocols
bgp {

```

```
group internal {
  type internal;
  local-address 198.51.100.0;
  cluster 198.51.100.1;
  neighbor 10.255.163.58 {
    description vpn1-to-pe1;
    family inet-vpn {
      unicast;
    }
    family route-target {
      proxy-generate;
    }
  }
  neighbor 10.255.168.42 {
    description vpn1-to-pe2;
    family inet-vpn {
      unicast;
    }
    family route-target;
  }
}
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@RR1# show routing-options
route-distinguisher-id 10.255.165.220;
autonomous-system 203;
```



## Configuring Device RR2

**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*.

To configure Device RR2:

1. Configure the interfaces.

```
[edit interfaces]
user@RR2# set ge-1/0/0 unit 0 description RR2-to-PE1
user@RR2# set ge-1/0/0 unit 0 family inet address 10.49.10.2/30
user@RR2# set ge-1/0/0 unit 0 family mpls

user@RR2# set ge-1/0/1 unit 0 description RR2-to-PE2
user@RR2# set ge-1/0/1 unit 0 family inet address 10.50.10.2/30
user@RR2# set ge-1/0/1 unit 0 family mpls
```

2. Configure the route distinguisher and the AS number.

```
[edit routing-options]
user@RR2# set route-distinguisher-id 10.255.165.28
user@RR2# set autonomous-system 203
```

3. Configure LDP as the signaling protocol used by the VPN.

```
[edit protocols ldp]
user@RR2# set interface ge-1/0/0
user@RR2# set interface ge-1/0/1
```

4. Configure BGP.

```
[edit protocols bgp group internal]
user@RR2# set type internal
user@RR2# set local-address 10.255.165.28
user@RR2# set cluster 198.51.100.1
user@RR2# set neighbor 10.255.163.58 description vpn2-to-pe1 family inet-vpn
unicast
user@RR2# set neighbor 10.255.168.42 description vpn2-to-pe2 family inet-vpn
unicast
```

5. Configure BGP route target filtering on the peering session with Device PE2.

```
[edit protocols bgp group internal]
user@RR2# set neighbor 10.255.168.42 family route-target
```

6. Configure proxy BGP route target filtering on the peering session with Device PE1.

```
[edit protocols bgp group internal]
user@RR2# set neighbor 10.255.163.58 family route-target proxy-generate
```

## 7. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@RR2# set interface ge-1/0/0
user@RR2# set interface ge-1/0/1
user@RR2# set interface lo0.0 passive
```

## 8. If you are done configuring the device, commit the configuration.

```
[edit]
user@RR2# commit
```

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

```
user@RR2# show interfaces
ge-1/0/0 {
  unit 0 {
    description RR2-to-PE1;
    family inet {
      address 10.49.10.2/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description RR2-to-PE2;
    family inet {
      address 10.50.10.2/30;
    }
    family mpls;
  }
}

user@RR2# show protocols
bgp {
  group internal {
    local-address 10.255.165.28;
    cluster 198.51.100.1;
    neighbor 10.255.163.58 {
      description vpn2-to-pe1;
      family inet-vpn {
        unicast;
      }
      family route-target {
        proxy-generate;
      }
    }
  }
  neighbor 10.255.168.42 {
    description vpn2-to-pe2;
    family inet-vpn {
      unicast;
    }
  }
}
```

```

    }
    family route-target;
  }
}
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@RR2# show routing-options
route-distinguisher-id 10.255.165.28;
autonomous-system 203;

```

### Configuring Device PE2

**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*.

To configure Device PE2:

1. Configure the interfaces.

```

[edit interfaces]
user@PE2# set ge-1/0/0 unit 0 description PE2-to-RR1
user@PE2# set ge-1/0/0 unit 0 family inet address 10.50.0.1/30
user@PE2# set ge-1/0/0 unit 0 family mpls

```

```

user@PE2# set ge-1/0/1 unit 0 description PE2-to-RR2
user@PE2# set ge-1/0/1 unit 0 family inet address 10.50.10.2/30
user@PE2# set ge-1/0/1 unit 0 family mpls

```

2. Configure the route distinguisher and the AS number.

```

[edit routing-options]
user@PE2# set route-distinguisher-id 10.255.168.42
user@PE2# set autonomous-system 203

```

3. Configure LDP as the signaling protocol used by the VPN.

```

[edit protocols ldp]
user@PE2# set interface ge-1/0/0
user@PE2# set interface ge-1/0/1

```

## 4. Configure BGP.

```
[edit protocols bgp group internal]
user@PE2# set type internal
user@PE2# set local-address 10.255.168.42
user@PE2# set family inet-vpn unicast
user@PE2# set family route-target
user@PE2# set neighbor 10.255.165.220
user@PE2# set neighbor 10.255.165.28
```

## 5. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@PE2# set interface ge-1/0/0
user@PE2# set interface ge-1/0/1
user@PE2# set interface lo0.0 passive
```

## 6. Configure the VPN routing instances.

```
[edit routing-instances vpn1]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:100
user@PE2# set routing-options static route 203.0.113.1/24 discard

[edit routing-instances vpn2]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:101
user@PE2# set routing-options static route 203.0.113.2/24 discard

[edit routing-instances vpn3]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:103
user@PE2# set routing-options static route 203.0.113.3/24 discard

[edit routing-instances vpn4]
user@PE2# set instance-type vrf
user@PE2# set vrf-target target:203:104
user@PE2# set routing-options static route 203.0.113.4/24 discard
```

## 7. Configure and apply the export routing policy.

```
[edit policy-options]
user@PE2# set rtf-prefix-list exclude-103 203:203:103/96

[edit policy-options policy-statement filter-rtc]
user@PE2# set from family route-target
user@PE2# set from rtf-prefix-list exclude-103
user@PE2# set then reject

[edit protocols bgp group internal]
user@PE2# set neighbor 10.255.165.220 export filter-rtc
```

## 8. If you are done configuring the device, commit the configuration.

```
[edit]
user@PE2# commit
```

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

```

user@PE2# show interfaces
ge-1/0/0 {
  unit 0 {
    description PE2-to-RR1;
    family inet {
      address 10.50.0.1/30;
    }
    family mpls;
  }
}
ge-1/0/1 {
  unit 0 {
    description PE2-to-RR2;
    family inet {
      address 10.50.10.2/30;
    }
    family mpls;
  }
}

user@PE2# show protocols
bgp {
  group internal {
    type internal;
    local-address 10.255.168.42;
    family inet-vpn {
      unicast;
    }
    family route-target;
    neighbor 10.255.165.220 {
      export filter-rtc;
    }
    neighbor 10.255.165.28;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-1/0/0.0;
    interface ge-1/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-1/0/0.0;
  interface ge-1/0/1.0;
}

user@PE2# show routing-options
route-distinguisher-id 10.255.168.42;

```

```
autonomous-system 203;

user@PE2# show policy-options
policy-statement filter-rtc {
  from {
    family route-target;
    rtf-prefix-list exclude-103;
  }
  then reject;
}
rtf-prefix-list exclude-103 {
  203:203:103/96;
}

user@PE2# show routing-instances
vpn1 {
  instance-type vrf;
  vrf-target target:203:100;
  routing-options {
    static {
      route 203.0.113.1/24 discard;
    }
  }
}
vpn2 {
  instance-type vrf;
  vrf-target target:203:101;
  routing-options {
    static {
      route 203.0.113.2/24 discard;
    }
  }
}
vpn3 {
  instance-type vrf;
  vrf-target target:203:103;
  routing-options {
    static {
      route 203.0.113.3/24 discard;
    }
  }
}
vpn4 {
  instance-type vrf;
  vrf-target target:203:104;
  routing-options {
    static {
      route 203.0.113.4/24 discard;
    }
  }
}
```

## Verification

Confirm that the configuration is working properly.

- Verifying the Route Target Filtering Routes in the `bgp.rtarget.0` Routing Table for Device RR1 on page 83
- Verifying the Route Target Filtering Routes in the `bgp.rtarget.0` Routing Table for Device RR2 on page 83

### Verifying the Route Target Filtering Routes in the `bgp.rtarget.0` Routing Table for Device RR1

**Purpose** Verify that the route prefix for `vpn3` is not in Device RR1's `bgp.rtarget.0` table. Since an export policy on Device PE2 was applied to prevent the advertisement of `vpn3` routes to Device RR1, Device RR1 should not receive those advertisements.

**Action** From operational mode, enter the `show route advertising-protocol bgp 10.255.165.220 table bgp.rtarget.0` command.

```
user@PE2# show route advertising-protocol bgp 10.255.165.220 table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 11 routes
(4 active, 0 holddown, 0 hidden)
  Prefix                Nexthop        MED      Lclpref    AS path
  203:203:100/96         *              Self      100         I
  203:203:101/96         *              Self      100         I
  203:203:104/96         *              Self      100         I
```

**Meaning** The `bgp.rtarget.0` table does not display `203:203:103/96`, which is the route prefix for `vpn3`. That means the export policy was applied correctly.

### Verifying the Route Target Filtering Routes in the `bgp.rtarget.0` Routing Table for Device RR2

**Purpose** Verify that the route prefix for `vpn3` is in Device RR2's `bgp.rtarget.0` table. Since an export policy was not applied on Device PE2 to prevent the advertisement of `vpn3` routes to Device RR2, Device RR2 should receive advertisements from all of the VPNs.

**Action** From operational mode, enter the `show route advertising-protocol bgp 10.255.165.28 table bgp.rtarget.0` command.

```
user@PE2# show route advertising-protocol bgp 10.255.165.28 table bgp.rtarget.0
bgp.rtarget.0: 4 destinations, 11 routes (4 active, 0 holddown, 0 hidden)
(4 active, 0 holddown, 0 hidden)
  Prefix                Nexthop        MED      Lclpref    AS path
  203:203:100/96         *              Self      100         I
  203:203:101/96         *              Self      100         I
  203:203:103/96         *              Self      100         I
  203:203:104/96         *              Self      100         I
```

**Meaning** The `bgp.rtarget.0` table displays the route prefixes for all of the VPNs.

- Related Documentation**
- [Example: Configuring Proxy BGP Route Target Filtering](#)
  - [Configuring BGP Route Target Filtering for VPNs on page 37](#)

---

## Reducing Network Resource Use with Static Route Target Filtering for VPNs

---

The BGP VPN route target extended community (RFC 4360, *BGP Extended Communities Attribute*) is used to determine VPN membership. Static route target filtering helps to prevent resources from being consumed in portions of the network where the VPN routes are not needed due to the lack of member PE routers (RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*). Routers can originate routes into the RT-Constrain protocol to indicate their interest in receiving VPN routes containing route targets that match the RT-Constrain NLRI.

Normally, for the RT-Constrain feature to function properly, it must be broadly deployed throughout a network. If this is not the case, the feature is less useful, because the RT-Constrain BGP speaker facing a non-RT-Constrain speaker must advertise a default RT-Constrain route to the other RT-Constrain speakers on behalf of the peer that does not support the feature. This effectively removes the resource saving benefits of the feature in portions of the network where it is not supported since a default RT-Constrain route causes the PE router and all intervening PE routers to need to receive all VPN routes.

The static RT-Constrain feature enables you to partially deploy the RT-Constrain feature in a network. The feature is enabled at a boundary in the network where RT-Constrain is configured. However, some BGP VPN peers do not support RT-Constrain, typically PE routers. The route targets of those PE routers must be statically configured on the router. These route targets are disseminated using the RT-Constrain protocol.

The proxy RT-Constrain feature permits BGP VPN peers that do not support the protocol to have their route-targets discovered and disseminated automatically. However, this feature can only support symmetric route-targets. For example, the import and export route-targets for a VRF routing instance are identical. However, for a hub-and-spoke VPN, the import and export route-targets are not identical. In this scenario, the import and export route-target may be statically configured to be disseminated in the RT-Constrain protocol.

- Related Documentation**
- [Configuring Static Route Target Filtering for VPNs on page 50](#)
  - [Configuring BGP Route Target Filtering for VPNs on page 37](#)



# Configuring Forwarding Options for VPNs

- [Chained Composite Next Hops for Layer 2 VPNs and Layer 2 Circuits on page 85](#)
- [Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs on page 86](#)

## Chained Composite Next Hops for Layer 2 VPNs and Layer 2 Circuits

The Juniper Networks PTX Series Packet Transport Routers and MX Series 3D Universal Edge Routers with MIC and MPC interfaces, and T4000 Core Routers are principally designed to handle large volumes of traffic in the core of large networks. Chained composite next hops help to facilitate this capability by allowing the router to process much larger volumes of routes. A chained composite next hop allows the router to direct sets of routes sharing the same destination to a common forwarding next hop, rather than having each route also include the destination. In the event that a network destination is changed, rather than having to update all of the routes sharing that destination with the new information, just the shared forwarding next hop is updated with the new information. The chained composite next hops continue to point to this forwarding next hop which now contains the new destination.

When the next hops for MPLS LSPs are created on the routers, the tag information corresponding to the inner-most MPLS label is extracted into a chained composite next hop. The chained composite next hop is stored in the ingress PFE. The chained composite next hop points to a next hop called the forwarding next hop that resides on the egress PFE. The forwarding next hop contains all of the other information (all of the labels except for the inner-most labels; and the IFA/IP information corresponding to the actual next hop node). Many chained composite next hops can share the same forwarding next hop. Additionally, separating the inner-most label (i.e the VPN label) from the forwarding next hop and storing it on the ingress PFE (within the chained composite next hop) helps to conserve egress PFE memory by reducing the number of rewrite strings stored on the egress PFE.

[Table 3 on page 85](#) shows support for chained composite next hops for ingress or transit routers on the MPLS network.

Table 3: Support for Chained Composite Next Hops

Platform	L2 VPN	L3 VPN	L2 CKT
PTX	Ingress and Transit	Ingress and Transit	Ingress only

Table 3: Support for Chained Composite Next Hops (continued)

Platform	L2 VPN	L3 VPN	L2 CKT
MX Series	Ingress only	Ingress only	Ingress only

To enable chained composite next hops on a T4000 router, the chassis must be configured to use the **enhanced-mode** option in network services mode.

For more information about configuring chassis network services, see the *Junos OS Administration Library*.

#### Related Documentation

- [Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928](#)
- [Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs on page 86](#)
- [chained-composite-next-hop on page 1045](#)
- [transit \(Chained Composite Next Hops\)](#)
- [ingress on page 1076](#)

## Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs

This example shows how to enable back-to-back Provider Edge (PE) router Layer 3 Virtual Private Network (VPN) connections with chained composite next hops for MIC and MPC interfaces on MX Series and T4000 routers.

- [Requirements on page 86](#)
- [Overview on page 87](#)
- [Configuration on page 88](#)
- [Verification on page 96](#)

### Requirements

This example uses the following hardware and software components:

- Six routers that can be a combination of MX240, MX480, MX960, or T4000 routers.
- Junos OS Release 13.3 running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all the routers:
  - a. MPLS
  - b. BGP

- c. LDP LSPs as tunnels between the PE devices
- d. OSPF or any other IGP protocol

## Overview

Prior to Junos OS Release 13.3, in a degenerated Layer 3 VPN case without the presence of an MPLS core router, previous behavior of flattened out indirect next hop and unicast next hop was utilized because there was no outer label available in the back-to-back PE-PE connection, and the ingress PE only pushed single VPN labels. In a Layer 3 VPN multipath scenario with mixed PE-PE and PE-P-PE paths, chained composite next hops could not be used either.

On platforms that support only MIC and MPC interfaces, chained composite next hops are enabled by default. On platforms that support both DPC and MPC interfaces, the Layer 3 VPN configuration required the **pe-pe-connection** statement to support chained composite next hops for PE-PE connections. However, the **pe-pe-connection** statement was not supported on platforms with MIC and FPC interfaces only.

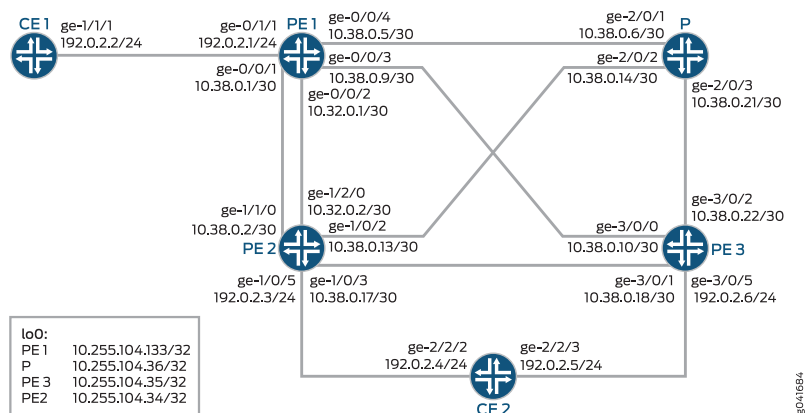
As a solution to these limitations, starting with Junos OS Release 13.3, the support for chained composite next hops is enhanced to automatically identify the underlying platform capability on composite next hops at startup time, without relying on user configuration, and to decide the next hop type (composite or indirect) to embed in the Layer 3 VPN label. This enhances the support for back-to-back PE-PE connections in Layer 3 VPN with composite next hops, and eliminates the need for the **pe-pe-connection** statement.

To enable chained composite next hops for directly connected PE devices, in addition to including the **l3vpn** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress]** hierarchy level, make the following changes:

- On MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs, chained composite next hops are disabled by default. To enable chained composite next hops on the MX240, MX480, and MX960, the chassis must be configured to use the **enhanced-ip** option in network services mode.
- On T4000 Core Routers containing MPC and FPCs, chained composite next hops are disabled by default. To enable chained composite next hops on a T4000 router, the chassis must be configured to use the **enhanced-mode** option in network services mode.

## Topology

Figure 7: Chained Composite Next Hop for PE-PE Connections



## 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.

```

CE1  set interfaces ge-1/1/1 unit 0 family inet address 192.0.2.2/24
    set interfaces ge-1/1/1 unit 0 family iso
    set interfaces ge-1/1/1 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 198.51.100.1/24
    set protocols bgp group PE type external
    set protocols bgp group PE peer-as 200
    set protocols bgp group PE neighbor 192.0.2.1
    set routing-options autonomous-system 100

PE1  set interfaces ge-0/0/1 unit 0 family inet address 10.38.0.1/30
    set interfaces ge-0/0/1 unit 0 family mpls
    set interfaces ge-0/0/2 unit 0 family inet address 10.38.0.5/30
    set interfaces ge-0/0/2 unit 0 family mpls
    set interfaces ge-0/0/3 unit 0 family inet address 10.38.0.9/30
    set interfaces ge-0/0/3 unit 0 family mpls
    set interfaces ge-0/0/4 unit 0 family inet address 10.32.0.1/30
    set interfaces ge-0/0/4 unit 0 family mpls
    set interfaces ge-0/1/1 unit 0 family inet address 192.0.2.1/24
    set interfaces ge-0/1/1 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.104.133/32
    set chassis network-services enhanced-ip
    set routing-options forwarding-table chained-composite-next-hop ingress l3vpn
    set routing-options autonomous-system 200
    set routing-options forwarding-table export lbpp
    set protocols mpls interface 10.38.0.1/30
    set protocols mpls interface 10.32.0.1/30

```

```

set protocols mpls interface 10.38.0.5/30
set protocols mpls interface 10.38.0.9/30
set protocols bgp group PEs type internal
set protocols bgp group PEs local-address 10.255.104.133
set protocols bgp group PEs family inet unicast
set protocols bgp group PEs family inet-vpn unicast
set protocols bgp group PEs neighbor 10.255.104.134 local-preference 200
set protocols bgp group PEs neighbor 10.255.104.135
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lbpp then load-balance per-packet
set routing-instances vpn-a instance-type vrf
set routing-instances vpn-a interface ge-0/1/1.0
set routing-instances vpn-a route-distinguisher 200:1
set routing-instances vpn-a vrf-target target:200:1
set routing-instances vpn-a vrf-table-label
set routing-instances vpn-a protocols bgp group CE type external
set routing-instances vpn-a protocols bgp group CE peer-as 100
set routing-instances vpn-a protocols bgp group CE neighbor 192.0.2.2

```

```

PE2 set interfaces ge-1/0/2 unit 0 family inet address 10.38.0.13/30
set interfaces ge-1/0/2 unit 0 family mpls
set interfaces ge-1/0/3 unit 0 family inet address 10.32.0.17/30
set interfaces ge-1/0/3 unit 0 family mpls
set interfaces ge-1/0/5 unit 0 family inet address 192.0.2.3/24
set interfaces ge-1/0/5 unit 0 family mpls
set interfaces ge-1/1/0 unit 0 family inet address 10.38.0.2/30
set interfaces ge-1/1/0 unit 0 family mpls
set interfaces ge-1/2/0 unit 0 family inet address 10.32.0.2/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.104.134/32
set chassis network-services enhanced-ip
set routing-options forwarding-table chained-composite-next-hop ingress l3vpn
set routing-instances vpn-a instance-type vrf
set routing-instances vpn-a interface ge-1/0/5.0
set routing-instances vpn-a route-distinguisher 200:2
set routing-instances vpn-a vrf-target target:200:1
set routing-instances vpn-a protocols bgp group CE type external
set routing-instances vpn-a protocols bgp group CE peer-as 300
set routing-instances vpn-a protocols bgp group CE neighbor 192.0.2.3
set protocols mpls interface 10.38.0.2/30
set protocols mpls interface 10.32.0.2/30
set protocols mpls interface 10.38.0.13/30
set protocols mpls interface 10.38.0.17/30
set protocols bgp group PEs type internal
set protocols bgp group PEs local-address 10.255.104.134
set protocols bgp group PEs family inet unicast
set protocols bgp group PEs family inet-vpn unicast
set protocols bgp group PEs neighbor 10.255.104.133
set protocols bgp group PEs neighbor 10.255.104.135
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

```

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-options autonomous-system 200

```

```

P    set interfaces ge-2/0/1 unit 0 family inet address 10.38.0.6/30
    set interfaces ge-2/0/1 unit 0 family mpls
    set interfaces ge-2/0/2 unit 0 family inet address 10.38.0.14/30
    set interfaces ge-2/0/2 unit 0 family mpls
    set interfaces ge-2/0/3 unit 0 family inet address 10.38.0.21/30
    set interfaces ge-2/0/3 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.104.136/32
    set protocols mpls interface 10.38.0.6/30
    set protocols mpls interface 10.38.0.14/30
    set protocols mpls interface 10.38.0.21/30
    set protocols bgp group PEs type internal
    set protocols bgp group PEs local-address 10.255.104.136
    set protocols bgp group PEs family inet unicast
    set protocols bgp group PEs family inet-vpn unicast
    set protocols bgp group PEs neighbor 10.255.104.133
    set protocols ospf area 0.0.0.0 interface all
    set protocols ospf area 0.0.0.0 interface fxp0.0 disable
    set protocols ospf area 0.0.0.0 interface lo0.0 passive
    set protocols ldp interface all
    set protocols ldp interface fxp0.0 disable
    set routing-options autonomous-system 200

```

```

PE3  set interfaces ge-3/0/0 unit 0 family inet address 10.38.0.10/30r0-r3
    set interfaces ge-3/0/0 unit 0 family mpls
    set interfaces ge-3/0/1 unit 0 family inet address 10.38.0.18/30r0-r1-2
    set interfaces ge-3/0/1 unit 0 family mpls
    set interfaces ge-3/0/2 unit 0 family inet address 10.38.0.22/30
    set interfaces ge-3/0/2 unit 0 family mpls
    set interfaces ge-3/0/5 unit 0 family inet address 192.0.2.6/24r0-r1-1
    set interfaces ge-3/0/5 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.104.135/32
    set chassis network-services enhanced-mode
    set routing-options forwarding-table chained-composite-next-hop ingress l3vpn
    set routing-options autonomous-system 200
    set routing-instances vpn-a instance-type vrf
    set routing-instances vpn-a interface ge-3/0/5.0
    set routing-instances vpn-a route-distinguisher 200:3
    set routing-instances vpn-a vrf-target target:200:1
    set routing-instances vpn-a protocols bgp group CE type external
    set routing-instances vpn-a protocols bgp group CE peer-as 300
    set routing-instances vpn-a protocols bgp group CE neighbor 192.0.2.5
    set protocols mpls interface 10.38.0.10/30
    set protocols mpls interface 10.38.0.18/30
    set protocols mpls interface 10.38.0.22/30
    set protocols bgp group PEs type internal
    set protocols bgp group PEs local-address 10.255.104.135
    set protocols bgp group PEs family inet unicast
    set protocols bgp group PEs family inet-vpn unicast
    set protocols bgp group PEs neighbor 10.255.104.133

```

```

set protocols bgp group PEs neighbor 10.255.104.134
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable

```

```

CE2 set interfaces ge-2/2/2 unit 0 family inet address 192.0.2.4/24
set interfaces ge-2/2/2 unit 0 family mpls
set interfaces ge-2/2/3 unit 0 family inet address 192.0.2.5/24
set interfaces ge-2/2/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.2/24
set protocols bgp group PE type external
set protocols bgp group PE metric-out 50
set protocols bgp group PE peer-as 200
set protocols bgp group PE export s2b
set protocols bgp group PE neighbor 192.0.2.4
set protocols bgp group PE neighbor 192.0.2.5
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 300

```

### Configuring Multipath Layer 3 VPN with Chained Composite Next Hop

#### 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*.

To configure basic Layer 3 VPN with chained composite next hop on the PE1 router:



**NOTE:** Repeat this procedure for the PE2 and PE3 routers in the MPLS domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure the interfaces on the PE1 router.

#### PE1 to CE1

```

[edit interfaces]
user@PE1 # set ge-0/1/1 unit 0 family inet address 192.0.2.1/24
user@PE1 # set ge-0/1/1 unit 0 family mpls

```

#### PE1 to PE2

```

[edit interfaces]
user@PE1 # set ge-0/0/1 unit 0 family inet address 10.38.0.1/30
user@PE1 # set ge-0/0/1 unit 0 family mpls

```

```

user@PE1 # set ge-0/0/2 unit 0 family inet address 10.38.0.5/30
user@PE1 # set ge-0/0/2 unit 0 family mpls

```

**PE1 to P**

```
[edit interfaces]
user@PE1 # set ge-0/0/4 unit 0 family inet address 10.32.0.1/30
user@PE1 # set ge-0/0/4 unit 0 family mpls
```

**PE1 to PE3**

```
[edit interfaces]
user@PE1 # set ge-0/0/3 unit 0 family inet address 10.38.0.9/30
user@PE1 # set ge-0/0/3 unit 0 family mpls
```

**Loopback interface**

```
[edit interfaces]
user@PE1 # set lo0 unit 0 family inet address 10.255.104.133/32
```

2. Enable enhanced IP mode on the PE1 chassis.

```
[edit chassis]
user@PE1# set network-services enhanced-ip
```

3. Enable chained composite next hop on the global Layer 3 VPN.

```
[edit routing-options]
user@PE1# set forwarding-table chained-composite-next-hop ingress l3vpn
```

4. Configure the autonomous system for PE1.

```
[edit routing-options]
user@PE1# set autonomous-system 200
```

5. Export the policy configured for load balancing.

```
[edit routing-options]
user@PE1# set forwarding-table export lbpp
```

6. Configure MPLS on the PE1 interfaces connecting to the P router and other PE routers.

```
[edit protocols]
user@PE1# set mpls interface 10.38.0.1/30
user@PE1# set mpls interface 10.32.0.1/30
user@PE1# set mpls interface 10.38.0.5/30
user@PE1# set mpls interface 10.38.0.9/30
```

7. Configure the IBGP group for PE1 to peer with the PE2 and PE3 routers.

```
[edit protocols]
user@PE1# set bgp group PEs type internal
user@PE1# set bgp group PEs local-address 10.255.104.133
user@PE1# set bgp group PEs family inet unicast
user@PE1# set bgp group PEs family inet-vpn unicast
user@PE1# set bgp group PEs neighbor 10.255.104.134 local-preference 200
user@PE1# set bgp group PEs neighbor 10.255.104.135
```



8. Configure OSPF with traffic engineering capability on all the interfaces of PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface lo0.0 passive
```

9. Configure LDP on all the interfaces of PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable
```

10. Configure a policy to load-balance traffic on a per packet basis.

```
[edit policy-options]
user@PE1# set policy-statement lbpp then load-balance per-packet
```

11. Configure a VRF routing instance on the CE1-facing interface of PE1.

```
[edit routing-instances]
user@PE1# set vpn-a instance-type vrf
user@PE1# set vpn-a interface ge-0/1/1.0
```

12. Configure the routing instance parameters.

```
[edit routing-instances]
user@PE1# set vpn-a route-distinguisher 200:1
user@PE1# set vpn-a vrf-target target:200:1
user@PE1# set vpn-a vrf-table-label
```

13. Configure an EBGP group for the routing instance, so PE1 can peer with CE1.

```
[edit routing-instances]
user@PE1# set vpn-a protocols bgp group CE type external
user@PE1# set vpn-a protocols bgp group CE peer-as 100
user@PE1# set vpn-a protocols bgp group CE neighbor 192.0.2.2
```

## Results

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

```
PE1 user@PE1# show chassis
network-services enhanced-ip;

user@PE1# show interfaces
ge-0/0/1 {
  unit 0 {
```

```
        family inet {
            address 10.38.0.1/30;
        }
        family mpls;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.38.0.5/30;
        }
        family mpls;
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 10.38.0.9/30;
        }
        family mpls;
    }
}
ge-0/0/4 {
    unit 0 {
        family inet {
            address 10.32.0.1/30;
        }
        family mpls;
    }
}
ge-0/1/1 {
    unit 0 {
        family inet {
            address 192.0.2.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.104.133/32;
        }
    }
}

user@PE1# show protocols
mpls {
    interface 10.38.0.1/30;
    interface 10.32.0.1/30;
    interface 10.38.0.5/30;
    interface 10.38.0.9/30;
}
bgp {
    group PEs {
        type internal;
```

```

        local-address 10.255.104.133;
        family inet {
            unicast;
        }
        family inet-vpn {
            unicast;
        }
        neighbor 10.255.104.134 {
            local-preference 200;
        }
        neighbor 10.255.104.135;
    }
}
ospf {
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

user@PE1# show routing-options
autonomous-system 200;
forwarding-table {
    export lbpp;
    chained-composite-next-hop {
        ingress {
            l3vpn;
        }
    }
}

user@PE1# show routing-instances
vpn-a {
    instance-type vrf;
    interface ge-0/1/1.0;
    route-distinguisher 200:1;
    vrf-target target:200:1;
    vrf-table-label;
    protocols {
        bgp {
            group CE {
                type external;
                peer-as 100;
                neighbor 192.0.2.2;
            }
        }
    }
}

```

```
    }  
  }  
  user@PE1# show policy-options  
  policy-statement lbpp {  
    then {  
      load-balance per-packet;  
    }  
  }
```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Routes on page 96](#)
- [Verifying Chained Next Hop on Direct PE-PE Connection on page 98](#)

### Verifying the Routes

---

**Purpose** Verify that the Layer 3 VPN prefixes toward PE1-PE2 point to composite next hops.

**Action** From operational mode, run the **show route 198.51.100.2 table vpn-a extensive** command.

```
user@PE1> show route 198.51.100.2 table vpn-a extensive
```

```
vpn-a.inet.0: 7 destinations, 10 routes (7 active, 0 holddown, 0 hidden)
198.51.100.2/24 (2 entries, 1 announced)
TSI:
KRT in-kerne 198.51.100.2/3 -> {composite(720)}
Page 0 idx 0, (group CE type External) Type 1 val 938eaa8 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: [200] 300 I
    Communities: target:200:1
Path 198.51.100.2 from 10.255.104.133 Vector len 4. Val: 0
  *BGP Preference: 170/-101
    Route Distinguisher: 200:2
    Next hop type: Indirect
    Address: 0x9391654
    Next-hop reference count: 12
    Source: 10.255.104.133
    Next hop type: Router, Next hop index: 1048580
    Next hop: 10.32.0.2 via ge-0/0/2.0
    Session Id: 0x1
    Next hop: 10.38.0.2 via ge-0/0/1.0, selected
    Session Id: 0x3
    Protocol next hop: 10.255.104.133
    Push 300192
    Composite next hop: 0x93918a4 718 INH Session ID: 0x9
    Indirect next hop: 0x941c000 1048581 INH Session ID: 0x9
    State: <Secondary Active Int Ext ProtectionCand>
    Local AS: 200 Peer AS: 200
    Age: 28 Metric: 50 Metric2: 1
    Validation State: unverified
    Task: BGP_203.0.113.1.133+57173
    Announcement bits (2): 0-KRT 1-BGP_RT_Background
    AS path: 300 I
    Communities: target:200:1
    Import Accepted
    VPN Label: 300192
    Localpref: 100
    Router ID: 10.255.104.133
    Primary Routing Table bgp.l3vpn.0
    Composite next hops: 1
      Protocol next hop: 10.255.104.133 Metric: 1
      Push 300192
      Composite next hop: 0x93918a4 718 INH Session ID: 0x9
      Indirect next hop: 0x941c000 1048581 INH Session ID:
0x9
        Indirect path forwarding next hops: 2
          Next hop type: Router
          Next hop: 10.32.0.2 via ge-1/0/0.0
          Session Id: 0x1
          Next hop: 10.38.0.2 via ge-1/1/2.0
          Session Id: 0x3
        10.255.104.133/32 Originating RIB: inet.3
          Metric: 1 Node path count: 1
          Forwarding nexthops: 2
            Nexthop: 10.32.0.2 via ge-0/0/2.0
      BGP Preference: 170/-101
```

```

Route Distinguisher: 200:3
Next hop type: Indirect
Address: 0x9391608
Next-hop reference count: 9
Source: 10.255.104.131
Next hop type: Router, Next hop index: 722
Next hop: 10.38.0.10 via ge-0/0/1.0, selected
Session Id: 0x4
Protocol next hop: 10.255.104.131
Push 299936
Composite next hop: 0x9391690 723 INH Session ID: 0xb
Indirect next hop: 0x941c0fc 1048583 INH Session ID: 0xb
State: <Secondary NotBest Int Ext ProtectionCand>
Inactive reason: Not Best in its group - Router ID
Local AS: 200 Peer AS: 200
Age: 28 Metric: 50 Metric2: 1
Validation State: unverified
Task: BGP_203.0.113.1.131+63797
AS path: 300 I
Communities: target:200:1
Import Accepted
VPN Label: 299936
Localpref: 100
Router ID: 10.255.104.131
Primary Routing Table bgp.l3vpn.0
Composite next hops: 1
    Protocol next hop: 10.255.104.131 Metric: 1
    Push 299936
    Composite next hop: 0x9391690 723 INH Session ID: 0xb
    Indirect next hop: 0x941c0fc 1048583 INH Session ID:
0xb
        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 10.38.0.10 via ge-1/0/2.0
            Session Id: 0x4
10.255.104.131/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 1
    Nexthop: 10.38.0.10 via ge-1/0/2.0

```

**Meaning** The PE2 router is the composite next hop for PE1 to reach CE2.

### Verifying Chained Next Hop on Direct PE-PE Connection

**Purpose** Verify that chained next hop is generated for direct PE-PE connection on CE1.

**Action** From operational mode, run the **ping** command.

```

user@CE1> ping 192.0.2.4
!!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

```

**Meaning** Chained composite next hop is enabled for the PE1 to PE2 connection.

**Related  
Documentation**

- *Chained Composite Next Hops for Transit Devices for VPNs*





## CHAPTER 6

# Configuring Graceful Restart for VPNs

- [VPN Graceful Restart on page 101](#)
- [Configuring Graceful Restart for VPNs on page 102](#)
- [Enabling Unicast Reverse-Path Forwarding Check for VPNs on page 104](#)
- [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 105](#)

## VPN Graceful Restart

---

VPN graceful restart allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. Without graceful restart, a control plane restart disrupts any VPN services provided by the router.

For VPN graceful restart to function properly, the following items need to be configured on the PE router:

- BGP graceful restart must be active on the PE-to-PE sessions carrying any service-signaling data in the session's network layer reachability information (NLRI).
- OSPF, IS-IS, LDP, and RSVP graceful restart must be active, because routes added by these protocols are used to resolve VPN NLRIs.
- For other protocols (static, Routing Information Protocol [RIP], and so on), graceful restart functionality must also be active when these protocols are run between the PE and CE routers. Layer 2 VPNs do not rely on this because protocols are not configured between the PE and CE routers.

In VPN graceful restart, a restarting router completes the following procedures:

- Waits for all the BGP NLRI information from other PE routers before it starts advertising routes to its CE routers.
- Waits for all protocols in all routing instances to converge (or finish graceful restart) before sending CE router information to the other PE routers.

- Waits for all routing instance information (whether it is local configuration or advertisements from a remote peer router) to be processed before sending it to the other PE routers.
- Preserves all forwarding state information in the MPLS routing tables until new labels and transit routes are allocated and then advertises them to other PE routers (and CE routers in carrier-of-carriers VPNs).

Graceful restart is supported on Layer 2 VPNs, Layer 3 VPNs, and virtual-router routing instances.

## Configuring Graceful Restart for VPNs

---

You can configure graceful restart to enable a router to pass through intermediate convergence states that are hidden from the rest of the network. Graceful restart allows a router whose VPN control plane is undergoing a restart (restarting router) to continue to forward traffic while recovering its state from neighboring routers (helper routers).

The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. When a restart event occurs and graceful restart is enabled, the restarting router can still forward traffic during the restart period, and convergence in the network is not disrupted. The helper routers hide the restart event from other devices not directly connected to the restarting router. In other words, the restart is not visible to the rest of the network, and the restarting router is not removed from the network topology.

Without graceful restart, a control plane restart disrupts any VPN services provided by the router. Graceful restart is supported on Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, and VPLS.

The graceful restart request occurs only if the following conditions are met:

- The network topology is stable.
- The neighbor or peer routers cooperate.
- The restarting router is not already cooperating with another restart already in progress.
- The grace period does not expire.

Before you begin:

- Configure the devices for network communication.
- Configure the device interfaces.

Graceful restart is disabled by default. To enable VPN graceful restart:

1. Configure graceful restart globally.

```
[edit routing-options]  
user@host# set graceful-restart
```

**NOTE:**

- Graceful restart can be enabled on logical systems. To configure graceful restart globally, include the `graceful-restart` statement at the [edit logical-systems *logical-system-name* routing-options] or the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options] hierarchy levels.
- To disable graceful restart globally, include the `disable` statement at the [edit routing-options graceful-restart] hierarchy level.

For example:

```
[edit routing-options]
user@host# set graceful-restart disable
```

2. Enable or disable graceful restart on a per-protocol, per-group, or per-neighbor basis, depending on the specific protocol, where the most specific definition is used.

```
[edit protocols]
user@host# set bgp graceful-restart
user@host# set bgp group group-name type internal local-address local-ip-address
neighbor neighbor1-address
user@host# set bgp group group-name type internal local-address local-ip-address
neighbor neighbor2-address graceful-restart disable
```

3. Configure graceful restart for Layer 3 VPNs for all routing and MPLS-related protocols within a routing instance. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported.

```
[edit routing-instance]
user@host# set routing-instance-name routing-options graceful-restart
```

**NOTE:**

- To disable graceful restart globally, include the `disable` statement at the [edit routing-instances *routing-instance-name* routing-options graceful-restart] hierarchy level.

For example:

```
[edit routing-instances]
user@host# set instance1 routing-options graceful-restart disable
```

- To disable graceful restart for individual protocols, include the `disable` statement at the [edit routing-instances *routing-instance-name* protocols *protocol-name* graceful-restart] hierarchy level.

For example:

```
[edit routing-instances]
user@host# set instance1 protocols ospf graceful-restart disable
```

4. Configure the duration of the graceful restart period for the routing instance.

```
[edit routing-options]
user@host# set graceful-restart restart-duration seconds
```

The **restart-duration** option sets the period of time that the router waits for a graceful restart to be completed. You can configure a time between 1 through 600 seconds. The default value is 300 seconds. At the end of the configured time period, the router performs a standard restart without recovering its state from the neighboring routers. This disrupts VPN services, but is probably necessary if the router is not functioning normally.



**NOTE:** You can include the **restart-duration** option at either the global or routing instance level. The routing instance value overrides the global value if both are configured.

#### Related Documentation

- *Graceful Restart and Layer 2 and Layer 3 VPNs*
- *Graceful Restart System Requirements*
- *Graceful Restart Concepts*
- *Verifying Graceful Restart Operation*

---

## Enabling Unicast Reverse-Path Forwarding Check for VPNs

IP spoofing may occur during a denial-of-service (DoS) attack. IP spoofing allows an intruder to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination. This type of spoofing is harmful because it consumes the destination's resources.

Unicast reverse-path forwarding (RPF) check is a tool to reduce forwarding of IP packets that may be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router forwards the packet to the destination address. If it is not from a valid path, the router discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family. You can also enable unicast RPF within a VPN routing instance.

To enable unicast RPF check, include the **unicast-reverse-path** statement:

```
unicast-reverse-path (active-paths | feasible-paths);
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To consider only active paths during the unicast RPF check, include the **active-paths** option. To consider all feasible paths during the unicast RPF check, include the **feasible-paths** option.

For more information about how to configure the **unicast-reverse-path** statement, see “[Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks](#)” on page 105 and .

- Related Documentation**
- [Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks on page 105](#)

## Example: Configuring Unicast Reverse-Path-Forwarding Check

- [Understanding How Unicast Reverse Path Forwarding Prevents Spoofed IP Packet Forwarding on page 105](#)
- [Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks on page 105](#)

## Understanding How Unicast Reverse Path Forwarding Prevents Spoofed IP Packet Forwarding

IP spoofing can occur during a denial-of-service (DoS) attack. IP spoofing allows an intruder to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination. This type of spoofing is harmful because it consumes the destination's resources.

A unicast reverse-path-forwarding (RPF) check is a tool to reduce forwarding of IP packets that might be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router or switch determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router or switch forwards the packet to the destination address. If it is not from a valid path, the router or switch discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family.



**NOTE:** Reverse path forwarding is not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.

- See Also**
- [Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks on page 105](#)

## Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks

Unicast reverse path forwarding (RPF) helps protect against DoS and DDoS attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast RPF to filter incoming traffic.

- [Requirements on page 106](#)
- [Overview on page 106](#)
- [Configuration on page 107](#)
- [Verification on page 112](#)

---

## Requirements

In this example, no special configuration beyond device initialization is required.

---

## Overview

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a DoS attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and DDoS attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding table entry for its source address. If the device uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the device forwards the packet. If the device does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the device discards the packet.


In this example, Device B has unicast RPF configured. Device A is using OSPF to advertise a prefix for the link that connects to Device D. OSPF is enabled on the links between Device B and Device C and the links between Device A and Device C, but not on the links between Device A and Device B. Therefore, Device B learns about the route to Device D through Device C.

If ingress filtering is used in an environment where DHCP or BOOTP is used, it should be ensured that the packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255 are allowed to reach the relay agent in routers when appropriate.

This example also includes a fail filter. When a packet fails the unicast RPF check, the fail filter is evaluated to determine if the packet should be accepted anyway. The fail filter in this example allows Device B's interfaces to accept Dynamic Host Configuration Protocol (DHCP) packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

[Figure 8 on page 107](#) shows the sample network.

Figure 8: Unicast RPF Sample Topology



### 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 A

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces fe-0/0/2 unit 5 family inet address 10.0.0.5/30
set interfaces fe-0/0/1 unit 17 family inet address 10.0.0.17/30
set interfaces fe-0/1/1 unit 25 family inet address 10.0.0.25/30
set interfaces fe-1/1/1 unit 29 family inet address 10.0.0.29/30
set protocols ospf export send-direct
set protocols ospf area 0.0.0.0 interface fe-0/1/1.25
set protocols ospf area 0.0.0.0 interface fe-1/1/1.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct from route-filter 10.0.0.16/30 exact
set policy-options policy-statement send-direct then accept
```

#### Device B

```
set interfaces fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/1/1 unit 6 family inet address 10.0.0.6/30
set interfaces fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/1 unit 9 family inet address 10.0.0.9/30
set interfaces fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/0 unit 13 family inet address 10.0.0.13/30
set protocols ospf area 0.0.0.0 interface fe-0/1/1.9
set protocols ospf area 0.0.0.0 interface fe-0/1/0.13
set routing-options forwarding-table unicast-reverse-path active-paths
set firewall filter rpf-special-case-dhcp term allow-dhcp from source-address 0.0.0.0/32
set firewall filter rpf-special-case-dhcp term allow-dhcp from destination-address 255.255.255.255/32
set firewall filter rpf-special-case-dhcp term allow-dhcp then count rpf-dhcp-traffic
set firewall filter rpf-special-case-dhcp term allow-dhcp then accept
set firewall filter rpf-special-case-dhcp term default then log
set firewall filter rpf-special-case-dhcp term default then reject
```

#### Device C

```
set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
```

```
set interfaces fe-0/0/2 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/0/2 unit 21 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 26 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 30 family inet address 10.0.0.30/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface fe-0/0/2.14
set protocols ospf area 0.0.0.0 interface fe-1/2/2.26
set protocols ospf area 0.0.0.0 interface fe-1/2/1.30
```

Device D      `set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30`

Device E      `set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30`

### *Configuring Device A*

**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*.

To configure Device A:

1.    Configure the interfaces.

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

user@A# set fe-0/0/2 unit 5 family inet address 10.0.0.5/30

user@A# set fe-0/0/1 unit 17 family inet address 10.0.0.17/30

user@A# set fe-0/1/1 unit 25 family inet address 10.0.0.25/30

user@A# set fe-1/1/1 unit 29 family inet address 10.0.0.29/30
```

2.    Configure OSPF.

```
[edit protocols ospf]
user@A# set export send-direct
user@A# set area 0.0.0.0 interface fe-0/1/1.25
user@A# set area 0.0.0.0 interface fe-1/1/1.29
```

3.    Configure the routing policy.

```
[edit policy-options policy-statement send-direct]
user@A# set from protocol direct
user@A# set from route-filter 10.0.0.16/30 exact
user@A# set then accept
```

4.    If you are done configuring Device A, commit the configuration.

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



**Configuring Device B**

**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*.

To configure Device B:

1. Configure the interfaces.

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

user@B# set fe-1/1/1 unit 6 family inet address 10.0.0.6/30

user@B# set fe-0/1/1 unit 9 family inet address 10.0.0.9/30

user@B# set fe-0/1/0 unit 13 family inet address 10.0.0.13/30
```

2. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@B# set interface fe-0/1/1.9
user@B# set interface fe-0/1/0.13
```

3. Configure unicast RPF, and apply the optional fail filter.

```
[edit interfaces]
user@B# set fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
```

4. (Optional) Configure the fail filter that gets evaluated if a packet fails the RPF check.

```
[edit firewall filter rpf-special-case-dhcp]
user@B# set term allow-dhcp from source-address 0.0.0.0/32
user@B# set term allow-dhcp from destination-address 255.255.255.255/32
user@B# set term allow-dhcp then count rpf-dhcp-traffic
user@B# set term allow-dhcp then accept
user@B# set term default then log
user@B# set term default then reject
```

5. (Optional) Configure only active paths to be considered in the RPF check.

This is the default behavior.

```
[edit routing-options forwarding-table]
user@B# set unicast-reverse-path active-paths
```

6. If you are done configuring Device B, commit the configuration.

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

### Results

Confirm your configuration by issuing the **show firewall**, **show interfaces**, **show protocols**, **show routing-options**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device A user@A# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
fe-0/0/2 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
fe-0/0/1 {
  unit 17 {
    family inet {
      address 10.0.0.17/30;
    }
  }
}
fe-0/1/1 {
  unit 25 {
    family inet {
      address 10.0.0.25/30;
    }
  }
}
fe-1/1/1 {
  unit 29 {
    family inet {
      address 10.0.0.29/30;
    }
  }
}

user@A# show protocols
ospf {
  export send-direct;
  area 0.0.0.0 {
    interface fe-0/1/1.25;
```

```

        interface fe-1/1/1.29;
    }
}

user@A# show policy-options
policy-statement send-direct {
    from {
        protocol direct;
        route-filter 10.0.0.16/30 exact;
    }
    then accept;
}

```

**Device B**

```

user@B# show firewall
filter rpf-special-case-dhcp {
    term allow-dhcp {
        from {
            source-address {
                0.0.0.0/32;
            }
            destination-address {
                255.255.255.255/32;
            }
        }
        then {
            count rpf-dhcp-traffic;
            accept;
        }
    }
    term default {
        then {
            log;
            reject;
        }
    }
}

user@B# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.2/30;
        }
    }
}
fe-1/1/1 {
    unit 6 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.6/30;
        }
    }
}
fe-0/1/1 {
    unit 9 {
        family inet {

```

```
        rpf-check fail-filter rpf-special-case-dhcp;  
        address 10.0.0.9/30;  
    }  
}  
fe-0/1/0 {  
    unit 13 {  
        family inet {  
            rpf-check fail-filter rpf-special-case-dhcp;  
            address 10.0.0.13/30;  
        }  
    }  
}
```

```
user@B# show protocols
```

```
ospf {  
    area 0.0.0.0 {  
        interface fe-0/1/1.9;  
        interface fe-0/1/0.13;  
    }  
}
```

```
user@B# show routing-options
```

```
forwarding-table {  
    unicast-reverse-path active-paths;  
}
```

Enter the configurations on Device C, Device D, and Device E, as shown in [“CLI Quick Configuration” on page 107](#).

---

### Verification

Confirm that the configuration is working properly.

- [Confirm That Unicast RPF Is Enabled on page 112](#)
- [Confirm That the Source Addresses Are Blocked on page 113](#)
- [Confirm That the Source Addresses Are Unblocked on page 113](#)

#### ***Confirm That Unicast RPF Is Enabled***

**Purpose** Make sure that the interfaces on Device B have unicast RPF enabled.

**Action** user@B> show interfaces fe-0/1/0.13 extensive

```

Logical interface fe-0/1/0.13 (Index 73) (SNMP ifIndex 553) (Generation 208)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Traffic statistics:
  Input bytes :          999390
  Output bytes :         1230122
  Input packets:          12563
  Output packets:         12613
Local statistics:
  Input bytes :          998994
  Output bytes :         1230122
  Input packets:          12563
  Output packets:         12613
Transit statistics:
  Input bytes :           396          0 bps
  Output bytes :           0          0 bps
  Input packets:           0          0 pps
  Output packets:          0          0 pps
Protocol inet, MTU: 1500, Generation: 289, Route table: 22
Flags: Sendbcst-pkt-to-re, uRPF
RPF Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.0.0.12/30, Local: 10.0.0.13, Broadcast: 10.0.0.15,
Generation: 241

```

**Meaning** The **uRPF** flag confirms that unicast RPF is enabled on this interface.

#### *Confirm That the Source Addresses Are Blocked*

**Purpose** Use the **ping** command to make sure that Device B blocks traffic from unexpected source addresses.

**Action** From Device A, ping Device B's interfaces, using 10.0.0.17 as the source address.

```

user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.6 (10.0.0.6): 56 data bytes
^C
--- 10.0.0.6 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

```

**Meaning** As expected, the ping operation fails.

#### *Confirm That the Source Addresses Are Unblocked*

**Purpose** Use the **ping** command to make sure that Device B does not block traffic when the RPF check is deactivated.

**Action** 1. Deactivate the RPF check on one of the interfaces.

2. Rerun the ping operation.

```
user@B> deactivate interfaces fe-1/1/1.6 family inet rpf-check

user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=63 time=1.263 ms
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.263/1.289/1.316/0.027 ms
```

**Meaning** As expected, the ping operation succeeds.

**See Also**

- [Understanding How Unicast Reverse Path Forwarding Prevents Spoofed IP Packet Forwarding on page 105](#)

**Related Documentation**

- *Example: Enabling Indirect Next Hops on the Packet Forwarding Engine*

## CHAPTER 7

# Configuring Class of Service for VPNs

- [VPNs and Class of Service on page 115](#)
- [Rewriting Class of Service Markers and VPNs on page 115](#)

## VPNs and Class of Service

---

You can configure Junos class-of-service (CoS) features to provide multiple classes of service for VPNs. The CoS features are supported on Layer2 VPNs, Layer 3 VPNs, and VPLS. On the router, you can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

VPNs use the standard CoS configuration.

### Related Documentation

- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

## Rewriting Class of Service Markers and VPNs

---

A marker reads the current forwarding class and loss priority information associated with a packet and finds the chosen code point from a table. It then writes the code point information into the packet header. Entries in a marker configuration represent the mapping of the current forwarding class into a new forwarding class, to be written into the header.

You define markers in the rewrite rules section of the class-of-service (CoS) configuration hierarchy and reference them in the logical interface configuration. You can configure different rewrite rules to handle VPN traffic and non-VPN traffic. The rewrite rule can be applied to MPLS and IPv4 packet headers simultaneously, making it possible to initialize MPLS experimental (EXP) and IP precedence bits at LSP ingress.

For a detailed example of how to configure rewrite rules for MPLS and IPv4 packets and for more information about how to configure statements at the **[edit class-of-service]** hierarchy level, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

- Related Documentation**
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*



## CHAPTER 8

# Pinging VPNs

- [Pinging VPNs, VPLS, and Layer 2 Circuits on page 117](#)
- [Setting the Forwarding Class of the Ping Packets on page 118](#)
- [Pinging a VPLS Routing Instance on page 118](#)
- [Pinging a Layer 2 VPN on page 119](#)
- [Pinging a Layer 3 VPN on page 119](#)
- [Pinging a Layer 2 Circuit on page 119](#)

### Pinging VPNs, VPLS, and Layer 2 Circuits

For testing purposes, you can ping Layer 2 VPNs, Layer 3 VPNs, and Layer 2 circuits by using the **ping mpls** command. The **ping mpls** command helps to verify that a VPN or circuit has been enabled and tests the integrity of the VPN or Layer 2 circuit connection between the PE routers. It does not test the connection between a PE router and a CE router. To ping a VPLS routing instance, you issue a **ping vpls instance** command (see [“Pinging a VPLS Routing Instance” on page 118](#)).

You issue the **ping mpls** command from the ingress PE router of the VPN or Layer 2 circuit to the egress PE router of the same VPN or Layer 2 circuit. When you execute the **ping mpls** command, echo requests are sent as MPLS packets.

The payload is a User Datagram Protocol (UDP) packet forwarded to the address **127.0.0.1**. The contents of this packet are defined in RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. The label and interface information for building and sending this information as an MPLS packet is the same as for standard VPN traffic, but the time-to-live (TTL) of the innermost label is set to 1.

When the echo request arrives at the egress PE router, the contents of the packet are checked, and then a reply that contains the correct return is sent by means of UDP. The PE router sending the echo request waits to receive an echo reply after a timeout of 2 seconds (you cannot configure this value).

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router (the router receiving the MPLS echo packets) to be able to ping the VPN or Layer 2 circuit. You must also configure the address **127.0.0.1/32** on the egress PE router's **lo0** interface. If this is not configured, the egress PE router does not have this forwarding entry and therefore simply drops the incoming MPLS pings.

The **ping mpls** command has the following limitations:

- You cannot ping an IPv6 destination prefix.
- You cannot ping a VPN or Layer 2 circuit from a router that is attempting a graceful restart.
- You cannot ping a VPN or Layer 2 circuit from a logical system.

You can also determine whether an LSP linking two PE routers in a VPN is up by pinging the end point address of the LSP. The command you use to ping an MPLS LSP end point is **ping mpls lsp-end-point address**. This command tells you what type of LSP (RSVP or LDP) terminates at the address specified and whether that LSP is up or down.

For a detailed description of this command, see the *Junos Routing Protocols and Policies Command Reference*.

---

## Setting the Forwarding Class of the Ping Packets

When you execute the **ping mpls** command, the ping packets forwarded to the destination include MPLS labels. It is possible to set the value of the forwarding class for these ping packets by using the **exp** option with the **ping mpls** command. For example, to set the forwarding class to 5 when pinging a Layer 3 VPN, issue the following command:

```
ping mpls l3vpn westcoast source 192.0.2.0 prefix 192.0.2.1 exp 5 count 20 detail
```

This command would make the router attempt to ping the Layer 3 VPN **westcoast** using ping packets with an EXP forwarding class of 5. The default forwarding class used for the **ping mpls** command packets is 7.

---

## Pinging a VPLS Routing Instance

The **ping vpls instance** command uses a different command structure and operates in a different fashion than the **ping mpls** command used for VPNs and Layer 2 circuits. The **ping vpls instance** command is only supported on MX Series routers, the M120 router, the M320 router, and the T1600 router.

To ping a VPLS routing instance, use the following command:

```
ping vpls instance instance-name destination-mac address source-ip address <count  
number> <data-plane-response> <detail> <learning-vlan-id number> <logical-system  
logical-system-name>
```

Pinging a VPLS routing instance requires using the **ping vpls instance** command with a combination of the routing instance name, the destination MAC address, and the source IP address (IP address of the outgoing interface).

When you run this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code; these packets are not counted in the received packets count. They are accounted for separately.

For more details, including argument descriptions and additional options, see [ping vpls instance](#).

## Pinging a Layer 2 VPN

To ping a Layer 2 VPN, use one of the following commands:

- **ping mpls l2vpn interface *interface-name***

You ping an interface configured for the Layer 2 VPN on the egress PE router.

- **ping mpls l2vpn instance *l2vpn-instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number***

You ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier to test the integrity of the Layer 2 VPN connection (specified by the identifiers) between the ingress and egress PE routers.

### Related Documentation

- *Example: Configuring MPLS-Based Layer 2 VPNs*

## Pinging a Layer 3 VPN

To ping a Layer 3 VPN, use the following command:

```
ping mpls l3vpn l3vpn-name prefix prefix <count count>
```

You ping a combination of an IPv4 destination prefix and a Layer 3 VPN name on the egress PE router to test the integrity of the VPN connection between the ingress and egress PE routers. The destination prefix corresponds to a prefix in the Layer 3 VPN. However, the ping tests only whether the prefix is present in a PE router's VRF table. It does not test the connection between a PE router and a CE router.

## Pinging a Layer 2 Circuit

To ping a Layer 2 circuit, use one of the following commands:

- **ping mpls l2circuit interface *interface-name***

You ping an interface configured for the Layer 2 circuit on the egress PE router.

- **ping mpls l2circuit virtual-circuit neighbor <*prefix*> <*virtual-circuit-id*>**

You ping a combination of the IPv4 prefix and the virtual circuit identifier on the egress PE router to test the integrity of the Layer 2 circuit between the ingress and egress PE routers.



## PART 2

# Common Configuration for Layer 3 VPNs

- [Layer 3 VPN Overview on page 123](#)
- [Assigning Routing Instances to Layer 3 VPNs on page 137](#)
- [Creating Unique VPN Routes Using VRF Tables on page 141](#)
- [Distributing Routes in Layer 3 VPNs on page 169](#)
- [Associating Interfaces with Layer 3 VPNs on page 211](#)
- [Carrying IPv4 Traffic Over Layer 3 VPNs on page 217](#)
- [Carrying IPv6 Traffic Over Layer 3 VPNs on page 227](#)
- [Configuring an AS for Layer 3 VPNs on page 241](#)
- [Limiting Routes Using Route Resolution and Route Filters on page 255](#)
- [Configuring Protocol Family Qualifiers for Layer 3 VPNs on page 263](#)
- [Configuring Class of Service for Layer 3 VPNs on page 267](#)
- [Enabling Internet Access for Layer 3 VPNs on page 271](#)



## CHAPTER 9

# Layer 3 VPN Overview

- [Understanding Layer 3 VPNs on page 123](#)
- [Understanding Layer 3 VPN Forwarding Through the Core on page 127](#)
- [Understanding Layer 3 VPN Attributes on page 129](#)
- [Supported Layer 3 VPN Standards on page 130](#)
- [Introduction to Configuring Layer 3 VPNs on page 130](#)
- [Layer 3 VPNs for IPv4 and IPv6 Overview on page 133](#)

### Understanding Layer 3 VPNs

---

Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. Instead of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks.

Layer 3 VPN operates at the Layer 3 level of the OSI model, the Network layer. A Layer 3 VPN is composed of a set of customer sites that are connected over a service provider's existing public Internet backbone. A peer-to-peer model is used to connect to the customer sites, where the service providers learn the customer routes on peering with the customers. The common routing information is shared across the provider's backbone using multiprotocol BGP, and the VPN traffic is forwarded to the customer sites using MPLS.

Junos OS supports Layer 3 VPNs based on RFC 4364. The RFC describes VPNs using MPLS tunnels for connectivity, BGP to distribute reachability information, and an IP backbone for transport. Service providers use their IP backbones to link a set of customer sites belonging to the same VPN.

- [Components of a Layer 3 VPN on page 124](#)
- [Layer 3 VPN Terminology on page 124](#)
- [Layer 3 VPN Architecture on page 125](#)

## Components of a Layer 3 VPN

There are three primary types of MPLS VPNs: Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs. All types of MPLS VPNs share certain components:

- CE devices—Customer Edge (CE) devices at the customer premises that connect to the provider's network. Some models call these Customer Premises Equipment (CPE) devices.
- Customer network—Customer sites with CE devices that belong to the VPN.
- Provider network—The service provider backbone network running the MPLS backbone.
- P devices—Provider (P) devices within the core of the provider's network. Provider devices are not connected to any device at a customer site and are part of the tunnel between pairs of PE devices. Provider devices support label-switched path (LSP) functionality as part of the tunnel support, but do not support VPN functionality.
- PE devices—Provider Edge (PE) devices within a service provider core network that connect directly to a CE device at the customer's site.
- MP-BGP—PE devices use MP-BGP to distribute customer routes to the proper PE devices across the MPLS backbone.

## Layer 3 VPN Terminology

VPNs use a distinct terminology to identify components of the network:

- IP routing table (also called the global routing table)—This table contains service provider routes not included in a VRF. Provider devices need this table to be able to reach each other, while the VRF table is needed to reach all customer devices on a particular VPN. For example, a PE router with Interface A to a CE router and Interface B to a backbone P router places the Interface A addresses in the VRF and the Interface B addresses in the global IP routing table.
- Route Distinguisher—A 64-bit value prepended to an IP address. This unique tag helps identify the different customers' routes as packets flow across the same service provider tunnel.

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

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



Each routing instance that you configure on a PE router must have a unique route distinguisher. There are two possible formats:

- **as-number:number**—Where, *as-number* is an autonomous system (AS) number (a 2-byte value) in the range 1 through 65,535, and *number* is any 4-byte value. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the ISP or the customer AS number.
- **ip-address:number**—Where, *ip-address* is an IP address (a 4-byte value), and *number* is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the router-id statement, which is a public IP address in your assigned prefix range.
- **Route Target (RT)**—A 64-bit value used to identify the final egress PE device for customer routes in a particular VRF to enable complex sharing of routes. The route target defines which route is part of a VPN. A unique route target helps distinguish between different VPN services on the same router. Each VPN also has a policy that defines how routes are imported into the VRF table on the router. A Layer 2 VPN is configured with import and export policies. A Layer 3 VPN uses a unique route target to distinguish between VPN routes. For example, the RT enables the sharing of routes in a shared service network to multiple customers. Each VPN route can have one or more RTs. A PE device handles RTs as extended BGP community values and uses the RTs to install customer routes.
- **VPN-IPv4 routes**—A route consisting of a 96-bit sequence composed of a 64-bit RD tag prepended to a 32-bit IPv4 address. The PE devices export the VPN-IPv4 routes in IBGP sessions to the other provider devices. These routes are exchanged across the MPLS backbone using iBGP. When the outbound PE device receives the route, it strips off the route distinguisher and advertises the route to the connected CE devices, typically through standard BGP IPv4 route advertisements.
- **VRF**—The virtual routing and forwarding (VRF) table distinguishes the routes for different customers, as well as customer routes from provider routes on the PE device. These routes can include overlapping private network address spaces, customer-specific public routes, and provider routes on a PE device useful to the customer.

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

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

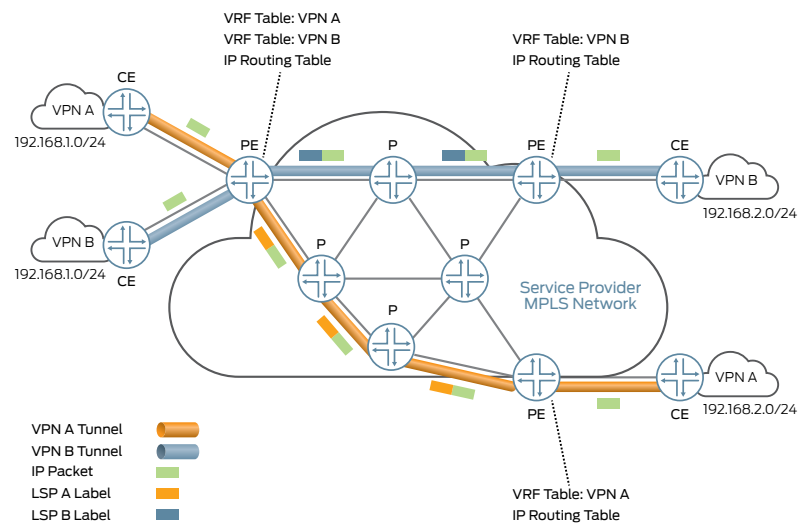
## Layer 3 VPN Architecture

A Layer 3 VPN links customer-edge routers (CE routers) to routers on the edge of the service provider network (PE routers). A Layer 3 VPN uses a peer routing model between local PE and CE routers that directly connect. That is, without needing multiple hops on the provider backbone to connect PE and CE router pairs. The PE routers distribute routing

information to all CE routers belonging to the same VPN, based on the BGP route distinguisher, locally and across the provider network. Each VPN has its own routing table for that VPN, coordinated with the routing tables in the CE and PE peer routers. The CE and PE routers have different VRF tables. Each CE router has only a single VRF table because the other VPNs are invisible to the CE. A PE router can connect to more than one CE router, so the PE router has a general IP routing table and VRF table for each attached CE with a VPN.

Figure 9 on page 126 shows the general architecture of a Layer 3 VPN.

**Figure 9: General Layer 3 VPN Architecture.**



The PE router knows which VRF table to use for packets arriving from remote VPN sites because every VRF table has one or more extended community attributes associated with it. The community attributes identify the route as belonging to a specific collection of routers. The route target community attribute identifies a collection of sites (more accurately, the collection of their VRF tables) to which a PE router distributes routes. The PE router uses the route target to import the correct remote VPN routes into its VRF tables.

The import and export of VPN routes between VPN sites is not automatic. This process is controlled by BGP routing policies. The routing policies establish the rules for exchanging routing information across the service provider's MPLS network and must be configured correctly and maintained when the network topology changes.

The PE router classifies IPv4 routes announced by a peer CE router and received by the PE router as VPN-IPv4 routes. When an ingress PE router receives routes advertised from a directly connected peer CE router, the ingress PE router checks the received route against the VRF export policy for that VPN. That is, the ingress PE router decides which remote PE routers need to know about the advertised routes. This is a two-step process:

- If the established export policy accepts the route, the PE router converts the information to VPN-IPv4 format by adding the route distinguisher to the IPv4 address. The PE router then announces the VPN-IPv4 route to the remote PE routers. The configured export

target policy of the VRF table determines the value of the attached route target. IBGP sessions distribute the VNP-IPv4 routes across the service provider's core network.

- If the established export policy does not accept the route, the PE router does not export the route to other PE routers, but the PE router uses the route locally. This happens, for example, when two CE routers in the same VPN connect directly to the same PE router so general traffic can flow from one CE site to another.

When an egress PE router on the other side of the service provider network receives a route, the egress PE router checks the route against the IBGP import policy in place between the PE routers. If the egress PE router accepts the route, then the egress PE router adds the route to the `bgp.l3vpn.0` routing table. The router also checks the route against the VRF import policy for the VPN. If the route is accepted, the egress PE router removes the route distinguisher and places the route into the correct VRF table. The VRF tables use the `routing-instance-name.inet.0` naming convention, so "VPN A" usually configures the table as `vpna.inet.0`.

#### Related Documentation

- [Understanding Virtual Routing and Forwarding Tables on page 141](#)
- [Understanding the BGP Local AS Attribute](#)
- [Understanding Layer 2 VPNs](#)
- [Interconnecting Layer 2 VPNs with Layer 3 VPNs Overview on page 957](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 935](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 936](#)
- [Example: Interconnecting a Layer 2 VPN with a Layer 3 VPN on page 958](#)

## Understanding Layer 3 VPN Forwarding Through the Core

The PE routers in the provider's core network are the only routers that are configured to support VPNs and hence are the only routers to have information about the VPNs. From the point of view of VPN functionality, the provider (P) routers in the core—those P routers that are not directly connected to CE routers—are merely routers along the tunnel between the ingress and egress PE routers.

The tunnels can be either LDP or MPLS. Any P routers along the tunnel must support the protocol used for the tunnel, either LDP or MPLS.

When PE-router-to-PE router forwarding is tunneled over MPLS label-switched paths (LSPs), the MPLS packets have a two-level label stack (see [Figure 10 on page 128](#)):

- Outer label—Label assigned to the address of the BGP next hop by the IGP next hop
- Inner label—Label that the BGP next hop assigned for the packet's destination address

Figure 10: Using MPLS LSPs to Tunnel Between PE Routers





Figure 11 on page 128 illustrates how the labels are assigned and removed:

1. When CE Router X forwards a packet to Router PE1 with a destination of CE Router Y, the PE route identifies the BGP next hop to Router Y and assigns a label that corresponds to the BGP next hop and identifies the destination CE router. This label is the inner label.
2. Router PE1 then identifies the IGP route to the BGP next hop and assigns a second label that corresponds to the LSP of the BGP next hop. This label is the outer label.
3. The inner label remains the same as the packet traverses the LSP tunnel. The outer label is swapped at each hop along the LSP and is then popped by the penultimate hop router (the third P router).
4. Router PE2 pops the inner label from the route and forwards the packet to Router Y.

Figure 11: Label Stack



## Understanding Layer 3 VPN Attributes

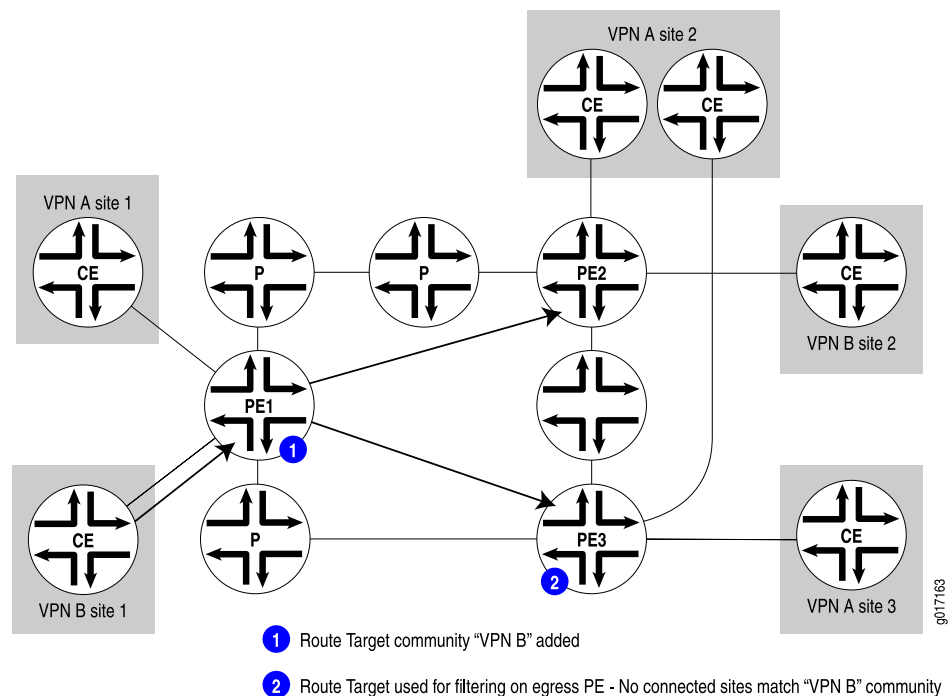
Route distribution within a VPN is controlled through BGP extended community attributes. RFC 4364 defines the following three attributes used by VPNs:

- **Target VPN**—Identifies a set of sites within a VPN to which a provider edge (PE) router distributes routes. This attribute is also called the *route target*. The route target is used by the egress PE router to determine whether a received route is destined for a VPN that the router services.

Figure 12 on page 129 illustrates the function of the route target. PE Router PE1 adds the route target “VPN B” to routes received from the customer edge (CE) router at Site 1 in VPN B. When it receives the route, the egress router PE2 examines the route target, determines that the route is for a VPN that it services, and accepts the route. When the egress router PE3 receives the same route, it does not accept the route because it does not service any CE routers in VPN B.

- **VPN of origin**—Identifies a set of sites and the corresponding route as having come from one of the sites in that set.
- **Site of origin**—Uniquely identifies the set of routes that a PE router learned from a particular site. This attribute ensures that a route learned from a particular site through a particular PE-CE connection is not distributed back to the site through a different PE-CE connection. It is particularly useful if you are using BGP as the routing protocol between the PE and CE routers and if different sites in the VPN have been assigned the same autonomous system (AS) numbers.

Figure 12: VPN Attributes and Route Distribution



## Supported Layer 3 VPN Standards

---

Junos OS substantially supports the following RFCs, which define standards for Layer 3 virtual private networks (VPNs).

- RFC 2283, *Multiprotocol Extensions for BGP-4*
- RFC 2685, *Virtual Private Networks Identifier*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*

The traceroute functionality is supported only on transit routers.

- RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*
- RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*

The following RFCs do not define a standard, but provide information about technology related to Layer 3 VPNs. The IETF classifies them as a “Best Current Practice” or “Informational.”

- RFC 1918, *Address Allocation for Private Internets*
- RFC 2917, *A Core MPLS IP VPN Architecture*

### Related Documentation

- [Supported Carrier-of-Carriers and Interprovider VPN Standards on page 313](#)
- [Supported VPWS Standards](#)
- [Supported Layer 2 VPN Standards](#)
- [Supported Multicast VPN Standards on page 473](#)
- [Supported VPLS Standards](#)
- [Supported MPLS Standards](#)
- [Supported Standards for BGP](#)
- [Accessing Standards Documents on the Internet](#)

## Introduction to Configuring Layer 3 VPNs

---

To configure Layer 3 virtual private network (VPN) functionality, you must enable VPN support on the provider edge (PE) router. You must also configure any provider (P) routers

that service the VPN, and you must configure the customer edge (CE) routers so that their routes are distributed into the VPN.

To configure Layer 3 VPNs, you include the following statements:

```
description text;
instance-type vrf;
interface interface-name;
protocols {
  bgp {
    group group-name {
      peer-as as-number;
      neighbor ip-address;
    }
    multihop tth-value;
  }
  (ospf | ospf3) {
    area area {
      interface interface-name;
    }
    domain-id domain-id;
    domain-vpn-tag number;
    sham-link {
      local address;
    }
    sham-link-remote address <metric number>;
  }
  rip {
    rip-configuration;
  }
}
route-distinguisher (as-number:id | ip-address:id);
router-id address;
routing-options {
  autonomous-system autonomous-system {
    independent-domain;
    loops number;
  }
  forwarding-table {
    export [ policy-names ];
  }
  interface-routes {
    rib-group group-name;
  }
  martians {
    destination-prefix match-type <allow>;
  }
  maximum-paths {
    path-limit;
    log-interval interval;
    log-only;
    threshold percentage;
  }
  maximum-prefixes {
    prefix-limit;
    log-interval interval;
  }
}
```

```

        log-only;
        threshold percentage;
    }
    multipath {
        vpn-unequal-cost;
    }
    options {
        syslog (level level | upto level);
    }
    rib routing-table-name {
        martians {
            destination-prefix match-type <allow>;
        }
        multipath {
            vpn-unequal-cost;
        }
        static {
            defaults {
                static-options;
            }
            route destination-prefix {
                next-hop [next-hops];
                static-options;
            }
        }
    }
}
static {
    defaults {
        static-options;
    }
    route destination-prefix {
        policy [ policy-names ];
        static-options;
    }
}
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-target (community | export community-name | import community-name);
vrf-table-label;

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]





**NOTE:** The `[edit logical-systems]` hierarchy level is not applicable in ACX Series routers.

The `sham-link`, `sham-link-remote`, and `vrf-advertise-selective` statements are not applicable in ACX Series routers.

For Layer 3 VPNs, only some of the statements in the `[edit routing-instances]` hierarchy are valid. For the full hierarchy, see *Junos OS Routing Protocols Library*.

In addition to these statements, you must enable a signaling protocol, IBGP sessions between the PE routers, and an interior gateway protocol (IGP) on the PE and P routers.

By default, Layer 3 VPNs are disabled.

Many of the configuration procedures for Layer 3 VPNs are common to all types of VPNs.

#### Related Documentation

- [Centralized Internet Access Through Layer 3 VPNs on page 296](#)
- [Configuring Hub-and-Spoke VPN Topologies: One Interface on page 653](#)
- [Configuring Overlapping VPNs Using Automatic Route Export on page 692](#)
- [Configuring a Full-Mesh VPN Topology with Route Reflectors on page 651](#)
- [Configuring a GRE Tunnel Interface Between a PE and CE Router on page 721](#)

## Layer 3 VPNs for IPv4 and IPv6 Overview

A Layer 3 virtual private network (VPN) routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces belong to the routing tables, and the routing protocol parameters control the information in the routing tables. In the case of MPLS VPNs, each VPN has a VPN routing and forwarding (VRF) routing instance.

A VRF routing instance consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table. Because each instance is configured for a particular VPN, each VPN has separate tables, rules, and policies that control its operation. A separate VRF table is created for each VPN that has a connection to a customer edge (CE) router. The VRF table is populated with routes received from directly connected CE sites associated with the VRF routing instance, and with routes received from other provider edge (PE) routers in the same VPN.

The standard or the global instance is called as the default routing instance. By default, all interfaces are associated with the default routing instance and default routing information base (RIB) (`inet0`). Routing options and routing policies supported on the default routing instance are also applicable to other routing instances.

A VRF routing instance is a BGP and MPLS VPN environment in which BGP is used to exchange IP VPN routes and discover the remote site, and in which VPN traffic traverses

an MPLS tunnel in an IP and MPLS backbone. You can enable an ACX Series router to function as a PE router by configuring VRF routing instances.

You can configure routing instances on ACX Series routers at the **[edit routing-instances routing-instance-name protocols]** hierarchy level for unicast IPv4, multicast IPv4, unicast IPv6, and multicast IPv6 address families. If you do not explicitly specify the address family in an IPv4 or an IPv6 environment, the router is configured to exchange unicast IPv4 or unicast IPv6 addresses by default. You can also configure the router to exchange unicast IPv4 and unicast IPv6 routes in a specified VRF routing instance. If you specify the multicast IPv4 or multicast IPv6 address family in the configuration, you can use BGP to exchange routing information about how packets reach a multicast source, instead of a unicast destination, for transmission to endpoints.



**NOTE:** Only the forwarding and virtual router routing instances support unicast IPv6 and multicast IPv6 address families. Unicast IPv6 and multicast IPv6 address families are not supported for VRF routing instances.

You can configure the following types of Layer 3 routing instances on ACX Series routers:

- **Forwarding**—Use this routing instance type for filter-based forwarding applications. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance **inet.0**. There are multiple forwarding tables and the selection of a table depends on the filter applied on the interface.
- **Virtual router**—A virtual router routing instance is similar to a VRF instance type, but is used for non-VPN-related applications. There are no VRF import, VRF export, VRF target, or route distinguisher requirements for this instance type. For this instance type, there is a one-to-one mapping between an interface and a routing instance. This routing instance type is used for routing and forwarding virtualization without VPNs (which is achieved by using the VRF-Lite application).
- **VRF**—Use the VRF routing instance type for Layer 3 VPN implementations. This routing instance type has a VPN routing table as well as a corresponding VPN forwarding table. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF routing instance corresponds with a forwarding table. The routes for each interface are installed in the forwarding table that is associated with the VRF routing instance. This routing instance type is used to implement BGP or MPLS VPNs in service provider networks or in big enterprise topologies.

Consider a sample VRF configuration scenario in which you want to configure two virtual routers, one to transmit voice and data traffic and another to carry management traffic. With such a configuration, the user and management networks are virtually separated, although the physical infrastructure is unified and cohesive. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your networks. The virtual routers do not create IP, MPLS, or GRE tunnels, and automatic discovery of remote sites that belong to the same network is not available. You must configure interfaces that are part of a virtual network in a streamlined manner to suit your topology requirements.

The following limitations apply to VRF routing instances that you configure on ACX Series routers:

- You cannot establish a communication between two virtual routing instances that are connected by external loopback.
- You cannot add a GRE or an MPLS tunnel to a virtual router.

In the Layer 3 lookup, up to 128 VRF tables are supported. Virtual routers without routing protocols enabled (based on static routes) support 64 VRF tables and virtual routers with all functions enabled within the routing instances support 16 VRF tables. When you enable VRF table labels and you do not explicitly apply a classifier configuration to the routing instance, the default MPLS EXP classifier is applied to the routing instance. You can override the default MPLS EXP classifier and apply a custom classifier to the routing instance. To perform this operation, you can filter the packets based on the IP header, choose the VRF, and based on the selected VRF, create an EXP classifier and associate it with the routing instance.

**Related  
Documentation**

- *Routing Instances Overview*
- [Configuring Virtual-Router Routing Instances in VPNs on page 14](#)
- *Applying MPLS EXP Classifiers to Routing Instances*
- [Configuring Routing Instances on PE Routers in VPNs on page 9](#)



# Assigning Routing Instances to Layer 3 VPNs

- [Understanding Routing Instances for Layer 3 VPNs on page 137](#)
- [Configuring Logical Units on the Loopback Interface for Routing Instances in Layer 3 VPNs on page 137](#)

## Understanding Routing Instances for Layer 3 VPNs

---

To implement Layer 3 VPNs in the JUNOS Software, you configure one routing instance for each VPN. You configure the routing instances on PE routers only. Each VPN routing instance consists of the following components:

- VRF table—On each PE router, you configure one VRF table for each VPN.
- Set of interfaces that use the VRF table—The logical interface to each directly connected CE router must be associated with a VRF table. You can associate more than one interface with the same VRF table if more than one CE router in a VPN is directly connected to the PE router.
- Policy rules—These control the import of routes into and the export of routes from the VRF table.
- One or more routing protocols that install routes from CE routers into the VRF table—You can use the BGP, OSPF, and RIP routing protocols, and you can use static routes.

## Configuring Logical Units on the Loopback Interface for Routing Instances in Layer 3 VPNs

---

For Layer 3 VPNs (VRF routing instances), you can configure a logical unit on the loopback interface into each VRF routing instance that you have configured on the router. Associating a VRF routing instance with a logical unit on the loopback interface allows you to easily identify the VRF routing instance.

Doing this is useful for troubleshooting:

- It allows you to ping a remote CE router from a local PE router in a Layer 3 VPN. For more information, see [“Example: Troubleshooting Layer 3 VPNs” on page 996](#).

- It ensures that a path maximum transmission unit (MTU) check on traffic originating on a VRF or virtual-router routing instance functions properly. For more information, see [“Configuring Path MTU Checks for VPN Routing Instances” on page 16](#).

You can also configure a firewall filter for the logical unit on the loopback interface; this configuration allows you to filter traffic for the VRF routing instance associated with it.

The following describes how firewall filters affect the VRF routing instance depending on whether they are configured on the default loopback interface, the VRF routing instance, or some combination of the two. The “default loopback interface” refers to **lo0.0** (associated with the default routing table), and the “VRF loopback interface” refers to **lo0.n**, which is configured in the VRF routing instance.

- If you configure Filter A on the default loopback interface and Filter B on the VRF loopback interface, the VRF routing instance uses Filter B.
- If you configure Filter A on the default loopback interface but do not configure a filter on the VRF loopback interface, the VRF routing instance does not use a filter.
- If you configure Filter A on the default loopback interface but do not even configure a VRF loopback interface, the VRF routing instance uses Filter A.

To configure a logical unit on the loopback interface, include the **unit** statement:

```
unit number {  
    family inet {  
        address address;  
    }  
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces lo0]**
- **[edit logical-systems *logical-system-name* interfaces lo0]**

To associate a firewall filter with the logical unit on the loopback interface, include the **filter** statement:

```
filter {  
    input filter-name;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces lo0 unit *unit-number* family inet]**
- **[edit logical-systems *logical-system-name* interfaces lo0 unit *unit-number* family inet]**

To include the **lo0.n** interface (where *n* specifies the logical unit) in the configuration for the VRF routing instance, include the following statement:

```
interface lo0.n;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

**Related  
Documentation**

- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*





# Creating Unique VPN Routes Using VRF Tables

- [Understanding Virtual Routing and Forwarding Tables on page 141](#)
- [Understanding VRF Localization in Layer 3 VPNs on page 145](#)
- [Maximizing VPN Routes Using VRF Localization for Layer 3 VPNs on page 146](#)
- [Example: Improving Scalability Using VRF Localization for Layer 3 VPNs on page 147](#)
- [Filtering Packets in Layer 3 VPNs Based on IP Headers on page 159](#)
- [Configuring a Label Allocation and Substitution Policy for VPNs on page 166](#)

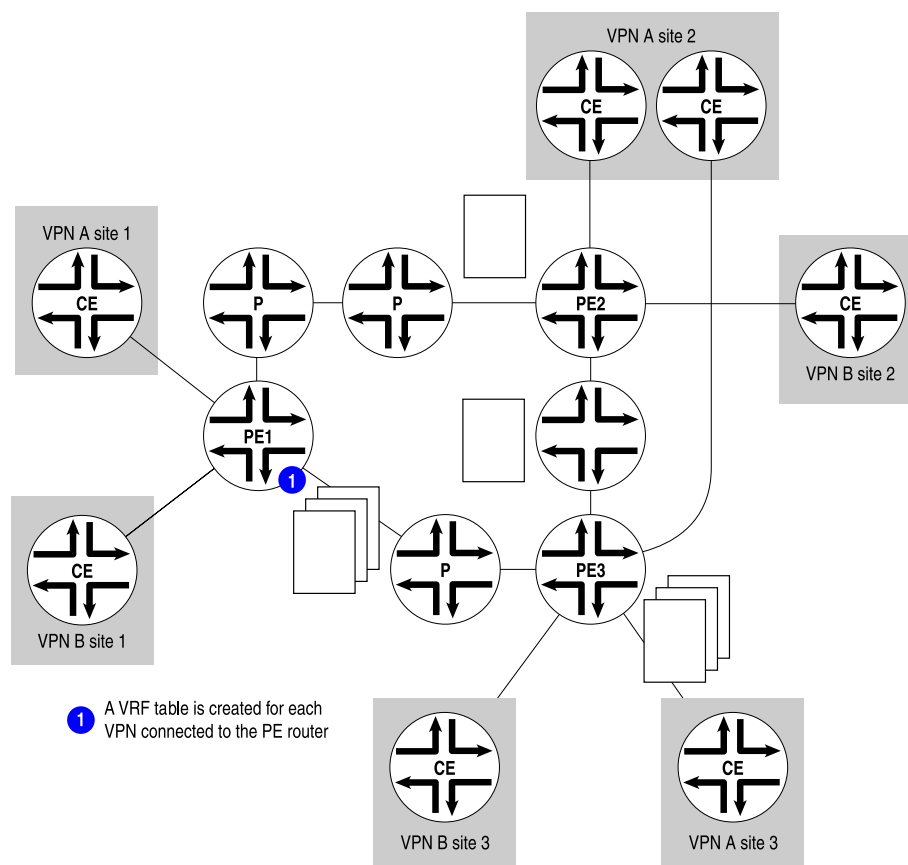
## Understanding Virtual Routing and Forwarding Tables

---

To separate a VPN's routes from routes in the public Internet or those in other VPNs, the PE router creates a separate routing table for each VPN, called a VPN routing and forwarding (VRF) table. The PE router creates one VRF table for each VPN that has a connection to a CE router. Any customer or site that belongs to the VPN can access only the routes in the VRF tables for that VPN.

[Figure 13 on page 142](#) illustrates the VRF tables that are created on the PE routers. The three PE routers have connections to CE routers that are in two different VPNs, so each PE router creates two VRF tables, one for each VPN.

Figure 13: VRF Tables



g017165

Each VRF table is populated from routes received from directly connected CE sites associated with that VRF routing instance and from routes received from other PE routers that passed BGP community filtering and are in the same VPN.

Each PE router also maintains one global routing table (`inet.0`) to reach other routers in and outside the provider's core network.

Each customer connection (that is, each logical interface) is associated with one VRF table. Only the VRF table associated with a customer site is consulted for packets from that site.

You can configure the router so that if a next hop to a destination is not found in the VRF table, the router performs a lookup in the global routing table, which is used for Internet access.

The Junos OS uses the following routing tables for VPNs:

- **bgp.l3vpn.0**—Stores all VPN-IPv4 unicast routes received from other PE routers. (This table does not store routes received from directly connected CE routers.) This table is present only on PE routers.

When a PE router receives a route from another PE router, it places the route into its **bgp.l3vpn.0** routing table. The route is resolved using the information in the **inet.3** routing table. The resultant route is converted into IPv4 format and redistributed to all **routing-instance-name.inet.0** routing tables on the PE router if it matches the VRF import policy.

The **bgp.l3vpn.0** table is also used to resolve routes over the MPLS tunnels that connect the PE routers. These routes are stored in the **inet.3** routing table. PE-to-PE router connectivity must exist in **inet.3** (not just in **inet.0**) for VPN routes to be resolved properly.

When a router is advertising non-local VPN-IPv4 unicast routes and the router is a route reflector or is performing external peering, the VPN-IPv4 unicast routes are automatically exported into the VPN routing table (**bgp.l3vpn.0**). This enables the router to perform path selection and advertise from the **bgp.l3vpn.0** routing table.

To determine whether to add a route to the **bgp.l3vpn.0** routing table, the Junos OS checks it against the VRF instance import policies for all the VPNs configured on the PE router. If the VPN-IPv4 route matches one of the policies, it is added to the **bgp.l3vpn.0** routing table. To display the routes in the **bgp.l3vpn.0** routing table, use the **show route table bgp.l3vpn.0** command.

- **routing-instance-name.inet.0**—Stores all unicast IPv4 routes received from directly connected CE routers in a routing instance (that is, in a single VPN) and all explicitly configured static routes in the routing instance. This is the VRF table and is present only on PE routers. For example, for a routing instance named **VPN-A**, the routing table for that instance is named **VPN-A.inet.0**.

When a CE router advertises to a PE router, the PE router places the route into the corresponding **routing-instance-name.inet.0** routing table and advertises the route to other PE routers if it passes a VRF export policy. Among other things, this policy tags the route with the route distinguisher (route target) that corresponds to the VPN site to which the CE belongs. A label is also allocated and distributed with the route. The **bgp.l3vpn.0** routing table is not involved in this process.

The **routing-instance-name.inet.0** table also stores routes announced by a remote PE router that match the VRF import policy for that VPN. The remote PE router redistributed these routes from its **bgp.l3vpn.0** table.

Routes are not redistributed from the **routing-instance-name.inet.0** table to the **bgp.l3vpn.0** table; they are directly advertised to other PE routers.

For each **routing-instance-name.inet.0** routing table, one forwarding table is maintained in the router's Packet Forwarding Engine. This table is maintained in addition to the forwarding tables that correspond to the router's **inet.0** and **mpls.0** routing tables. As with the **inet.0** and **mpls.0** routing tables, the best routes from the **routing-instance-name.inet.0** routing table are placed into the forwarding table.

To display the routes in the ***routing-instance-name.inet.0*** table, use the **show route table *routing-instance-name.inet.0*** command.

- **inet.3**—Stores all MPLS routes learned from LDP and RSVP signaling done for VPN traffic. The routing table stores the MPLS routes only if the **traffic-engineering bgp-igp** option is not enabled.

For VPN routes to be resolved properly, the **inet.3** table must contain routes to all the PE routers in the VPN.

To display the routes in the **inet.3** table, use the **show route table inet.3** command.

- **inet.0**—Stores routes learned by the IBGP sessions between the PE routers. To provide Internet access to the VPN sites, configure the ***routing-instance-name.inet.0*** routing table to contain a default route to the **inet.0** routing table.

To display the routes in the **inet.0** table, use the **show route table inet.0** command.

The following routing policies, which are defined in VRF import and export statements, are specific to VRF tables.

- Import policy—Applied to VPN-IPv4 routes learned from another PE router to determine whether the route should be added to the PE router's **bgp.l3vpn.0** routing table. Each routing instance on a PE router has a VRF import policy.
- Export policy—Applied to VPN-IPv4 routes that are announced to other PE routers. The VPN-IPv4 routes are IPv4 routes that have been announced by locally connected CE routers.

VPN route processing differs from normal BGP route processing in one way. In BGP, routes are accepted if they are not explicitly rejected by import policy. However, because many more VPN routes are expected, the Junos OS does not accept (and hence store) VPN routes unless the route matches at least one VRF import policy. If no VRF import policy explicitly accepts the route, it is discarded and not even stored in the **bgp.l3vpn.0** table. As a result, if a VPN change occurs on a PE router—such as adding a new VRF table or changing a VRF import policy—the PE router sends a BGP route refresh message to the other PE routers (or to the route reflector if this is part of the VPN topology) to retrieve all VPN routes so they can be reevaluated to determine whether they should be kept or discarded.

#### Related Documentation

- *IGP Shortcuts and VPNs*

## Understanding VRF Localization in Layer 3 VPNs

In a Layer 3 VPN, to separate routes of a VPN from routes in the public Internet or those in other VPNs, the PE router creates a separate routing table for each VPN, called a virtual routing and forwarding (VRF) table. Each VRF uses a route distinguisher and route target to differentiate other VPNs so that each VRF achieves a VPN in a public network. The PE router creates one VRF table for each VPN that has a connection to a CE router. Any customer or site that belongs to the VPN can access only the routes in the VRF tables for that VPN.

The PE routers in a Layer 3 VPN deployment have two types of line cards hosting the following interfaces:

- CE-facing interfaces
- Core-facing interfaces



**NOTE:** An FPC can be either core-facing or CE-facing.

The VRFs are present on these line cards and currently, in Junos OS, all the routes of all the VRFs are present on all line cards along with chained composite next hops on all the FPCs. This uses up the memory in each line card. Since traffic from CE-facing interfaces comes in only through the corresponding CE-facing FPCs, all the routes and next hops need not be present on all the line cards. VRF localization provides a mechanism for localizing routes of VRF to specific line cards to help maximize the number of routes that a router can handle. CE-facing interfaces localize all the routes of instance type VRF to a specific line card. If CE-facing interfaces are logical interfaces like AE or RLSQ or IRB, then a line card number has to be configured to localize routes. Core-facing line cards store all the VRF routes. These cards have to be configured as VPN core-facing default or VPN core-facing only. Core-facing line cards store routes of all the VRFs, and they are of the following types:

- vpn-core-facing-default — The core-facing FPC installs all the routes and next hops of the VRF routes.
- vpn-core-facing-only — The core-facing FPC installs all the routes and does not store next hops of the VRF routes.



**NOTE:** Core-facing FPCs can be configured as either core-facing-default or core-facing-only.

### Related Documentation

- [Example: Improving Scalability Using VRF Localization for Layer 3 VPNs on page 147](#)
- [localized-fib on page 1095](#)
- [vpn-localization on page 1150](#)

## Maximizing VPN Routes Using VRF Localization for Layer 3 VPNs

Virtual routing and forwarding (VRF) localization provides a mechanism for localizing routes of VRF to specific line cards to help maximize the number of routes that a router can handle. CE-facing interfaces localize all the routes of instance type VRF to a specific line card. If the CE-facing interfaces are logical interfaces like AE/RLSQ/IRB, then the line card has to be configured to localize routes. Core-facing line cards store all the VRF routes. These cards have to be configured as VPN core-facing only or VPN core-facing default. To configure VRF localization, configure the **localized-fib** statement at the **[edit routing-instances *instance-name* routing-options]** hierarchy level and configure the **vpn-localization** statement at the **[edit chassis fpc *fpc-slot*]** hierarchy level. The **show route vpn-localization** command displays the localization information of all the VRFs in the system.

Before you begin to localize the VRF table:

- Configure the interfaces.
- Configure the routing and signaling protocols.

To configure VRF localization:

1. Configure the chassis of the router.
  - a. Configure the FPC slot as either VPN core-facing only or VPN core-facing default to store the VRF routes.

```
[edit chassis]
user@host# set fpc slot-number vpn-localization vpn-core-facing-only
user@host# set fpc slot-number vpn-localization vpn-core-facing-default
```

2. Configure enhanced IP network service on the chassis.

```
[edit chassis]
user@host# set network-services enhanced-ip
```

3. Create an instance type, configure the route distinguisher, and configure the VRF target community and VRF target label.

```
[edit routing-instances routing-instance]
user@host# set instance-type vrf
user@host# set interface interface-name
user@host# set route-distinguisher route-distinguisher-id
user@host# set provider-tunnel rsvp-te static-lsp vpn1-p2mp
user@host# set vrf-target vrf-target-community
user@host# set vrf-table-label
```

4. Configure the multipath routing option to balance load independent of the protocol.

```
[edit routing-instances routing-instance routing-options]
user@host# set multipath
```

5. Configure the specific FPC of CE-facing physical interfaces or specify the FPC slot number if the CE-facing interfaces are logical interfaces like AE or RSQ or IRB to localize the VRF routing instance routes.

- Configure the specific FPC of CE-facing physical interfaces to localize the VRF routing instance routes.

```
[edit routing-instances routing-instance routing-options]
user@host# set localized-fib
```

- Configure the FPC slot number of the CE-facing logical interfaces like AE or RSQ or IRB to localize the VRF routing instance routes.

```
[edit routing-instances routing-instance routing-options]
user@host# set localized-fib fpc-slot fpc-slot-number
```

6. Configure the peer group of the BGP protocol for the routing instance.

```
[edit routing-instances routing-instance protocols bgp group group-name]
user@host# set type external
user@host# set export direct
user@host# set peer-as 100
user@host# set neighbor IP-address family inet unicast
user@host# set neighbor IP-address family inet6 unicast
```

7. Configure the MVPN protocol for the routing instance.

```
[edit routing-instances routing-instance protocols]
user@host# set mvpn
```

#### Related Documentation

- [Example: Improving Scalability Using VRF Localization for Layer 3 VPNs on page 147](#)
- [Understanding VRF Localization in Layer 3 VPNs on page 145](#)

## Example: Improving Scalability Using VRF Localization for Layer 3 VPNs

This example shows how to configure VRF localization on MX Series routers, which enables you to improve the VPN scalability on MX Series routers.

- [Requirements on page 147](#)
- [Overview on page 148](#)
- [Configuration on page 148](#)
- [Verification on page 158](#)

### Requirements

This example uses the following hardware and software components:

- Five MX Series 3D Universal Edge Routers
- Junos OS Release 14.2 or later running on all devices

Before you begin:

1. Configure the device interfaces.
2. Configure the BGP protocol.


## Overview

Starting with Junos OS Release 14.2, the VRF localization provides a mechanism for localizing routes of VRF to specific line cards which helps maximize the number of routes that a router can handle. CE-facing interfaces localize all the routes of instance type VRF to a specific line card. If the CE-facing interfaces are logical interfaces like AE or RLSQ or IRB, then the line card has to be configured to localize routes. Core-facing line cards store all the VRF routes. These cards have to be configured as VPN core-facing only or VPN core-facing default. To configure VRF localization, configure the **localized-fib** configuration statement at the **[edit routing-instances instance-name routing-options]** hierarchy level and configure **vpn-localization** at the **[edit chassis fpc fpc-slot]** hierarchy level. The **show route vpn-localization** command displays the localization information of all the VRFs in the system.

## Topology

In the topology shown in [Figure 14 on page 148](#), VRF localization is configured on Device PE1.

Figure 14: Example VRF Localization



## 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 the configuration mode.

```
CE1 set interfaces ge-4/0/0 unit 0 family inet address 192.0.2.2/24
    set interfaces ge-4/0/0 unit 0 family inet6 address abcd:a:a:a:1::2/126
    set protocols bgp group vpn1 type external
    set protocols bgp group vpn1 export direct
    set protocols bgp group vpn1 peer-as 10
    set protocols bgp group vpn1 neighbor 192.0.2.1 family inet unicast
    set protocols bgp group vpn1 neighbor abcd:a:a:a:1::1 family inet6 unicast
    set policy-options policy-statement direct from protocol direct
    set policy-options policy-statement direct then accept
    set policy-options policy-statement load-balancing-policy then load-balance per-packet
    set routing-options autonomous-system 100
    set routing-options forwarding-table export load-balancing-policy
```



```

PE1 set chassis redundancy graceful-switchover
    set chassis aggregated-devices ethernet device-count 16
    set chassis fpc 8 vpn-localization vpn-core-facing-only
    set chassis network-services enhanced-ip
    set interfaces ge-2/0/0 unit 0 family inet address 192.0.2.1/24
    set interfaces ge-2/0/0 unit 0 family inet6 address abcd:a:a:1::1/126
    set interfaces ge-8/1/0 gigether-options 802.3ad ae0
    set interfaces ge-8/1/9 gigether-options 802.3ad ae0
    set interfaces ae0 unit 0 family inet address 192.0.2.3/24
    set interfaces ae0 unit 0 family iso
    set interfaces ae0 unit 0 family mpls
    set interfaces lo0 unit 1 family inet address 203.0.113.1/24
    set interfaces lo0 unit 1 family inet6 address abcd::10:0:1:1/128
    set policy-options policy-statement direct from protocol direct
    set policy-options policy-statement direct then accept
    set policy-options policy-statement load-balancing-policy then load-balance per-packet
    set protocols rsvp interface ae0.0
    set protocols mpls ipv6-tunneling
    set protocols mpls icmp-tunneling
    set protocols mpls label-switched-path pe1-pe2-p2mp-1 from 10.255.19.254
    set protocols mpls label-switched-path pe1-pe2-p2mp-1 to 10.255.19.251
    set protocols mpls label-switched-path pe1-pe2-p2mp-1 link-protection
    set protocols mpls label-switched-path pe1-pe2-p2mp-1 p2mp vpn1-p2mp
    set protocols mpls label-switched-path pe1-pe3-p2mp-1 from 10.255.19.254
    set protocols mpls label-switched-path pe1-pe3-p2mp-1 to 10.255.19.203
    set protocols mpls label-switched-path pe1-pe3-p2mp-1 link-protection
    set protocols mpls label-switched-path pe1-pe3-p2mp-1 p2mp vpn1-p2mp
    set protocols mpls interface ae0.0
    set protocols bgp group mpbg type internal
    set protocols bgp group mpbg local-address 10.255.19.254
    set protocols bgp group mpbg family inet unicast
    set protocols bgp group mpbg family inet-vpn unicast
    set protocols bgp group mpbg family inet6 unicast
    set protocols bgp group mpbg family inet6-vpn unicast
    set protocols bgp group mpbg family inet-mvpn signaling
    set protocols bgp group mpbg family inet6-mvpn signaling
    set protocols bgp group mpbg neighbor 10.255.19.253
    set protocols ospf traffic-engineering
    set protocols ospf area 0.0.0.0 interface ae0.0
    set protocols ospf area 0.0.0.0 interface lo0.0 passive
    set protocols ldp interface ae0.0
    set routing-instances vpn1 instance-type vrf
    set routing-instances vpn1 interface ge-2/0/0.0
    set routing-instances vpn1 interface lo0.1
    set routing-instances vpn1 route-distinguisher 1:1
    set routing-instances vpn1 provider-tunnel rsvp-te static-lsp vpn1-p2mp
    set routing-instances vpn1 vrf-target target:1:1
    set routing-instances vpn1 vrf-table-label
    set routing-instances vpn1 routing-options multipath
    set routing-instances vpn1 routing-options localized-fib
    set routing-instances vpn1 protocols bgp group grp1 type external
    set routing-instances vpn1 protocols bgp group grp1 export direct
    set routing-instances vpn1 protocols bgp group grp1 peer-as 100
    set routing-instances vpn1 protocols bgp group grp1 neighbor 192.0.2.2 family inet unicast
    set routing-instances vpn1 protocols bgp group grp1 neighbor abcd:a:a:1::2 family inet6
    unicast

```

```

set routing-instances vpn1 protocols mvpn
set routing-options nonstop-routing
set routing-options autonomous-system 10
set routing-options forwarding-table export load-balancing-policy
set routing-options forwarding-table chained-composite-next-hop ingress l3vpn
  extended-space

```

```

P    set chassis aggregated-devices ethernet device-count 16
      set interfaces ge-1/0/1 gigether-options 802.3ad ae0
      set interfaces ge-1/0/3 gigether-options 802.3ad ae0
      set interfaces ge-1/1/1 gigether-options 802.3ad ae1
      set interfaces ae0 unit 0 family inet address 192.0.2.4/24
      set interfaces ae0 unit 0 family iso
      set interfaces ae0 unit 0 family mpls
      set interfaces ae1 unit 0 family inet address 198.51.100.2/24
      set interfaces ae1 unit 0 family iso
      set interfaces ae1 unit 0 family mpls
      set routing-options autonomous-system 10
      set routing-options forwarding-table export load-balancing-policy
      set protocols rsvp interface ae0.0
      set protocols rsvp interface ae1.0
      set protocols mpls ipv6-tunneling
      set protocols mpls icmp-tunneling
      set protocols mpls interface ae0.0
      set protocols mpls interface ae1.0
      set protocols bgp group mpbg type internal
      set protocols bgp group mpbg local-address 10.255.19.253
      set protocols bgp group mpbg family inet unicast
      set protocols bgp group mpbg family inet-vpn unicast
      set protocols bgp group mpbg family inet6 unicast
      set protocols bgp group mpbg family inet6-vpn unicast
      set protocols bgp group mpbg family inet-mvpn signaling
      set protocols bgp group mpbg family inet6-mvpn signaling
      set protocols bgp group mpbg cluster 10.255.19.253
      set protocols bgp group mpbg neighbor 10.255.19.254
      set protocols bgp group mpbg neighbor 10.255.19.251
      set protocols bgp group mpbg neighbor 10.255.19.203
      set protocols ospf traffic-engineering
      set protocols ospf area 0.0.0.0 interface lo0.0 passive
      set protocols ospf area 0.0.0.0 interface ae0.0
      set protocols ospf area 0.0.0.0 interface ae1.0
      set protocols ldp interface ae0.0
      set protocols ldp interface ae1.0
      set policy-options policy-statement load-balancing-policy then load-balance per-packet

```

```

PE2  set chassis redundancy graceful-switchover
      set chassis aggregated-devices ethernet device-count 16
      set interfaces ge-4/2/1 gigether-options 802.3ad ae1
      set interfaces ge-4/2/5 unit 0 family inet address 198.51.100.3/24
      set interfaces ge-4/2/5 unit 0 family inet6 address abcd:a:a:a:2::1/126
      set interfaces ae1 unit 0 family inet address 198.51.100.1/24
      set interfaces ae1 unit 0 family iso
      set interfaces ae1 unit 0 family mpls
      set interfaces lo0 unit 2 family inet address 203.0.113.2/24

```

```

set interfaces lo0 unit 2 family inet6 address abcd::203:0:113:2/128
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set protocols rsvp interface ae1.0
set protocols mpls ipv6-tunneling
set protocols mpls icmp-tunneling
set protocols mpls label-switched-path pe2-pe1-p2mp-1 from 10.255.19.251
set protocols mpls label-switched-path pe2-pe1-p2mp-1 to 10.255.19.254
set protocols mpls label-switched-path pe2-pe1-p2mp-1 link-protection
set protocols mpls label-switched-path pe2-pe1-p2mp-1 p2mp vpn1-p2mp
set protocols mpls label-switched-path pe2-pe3-p2mp-1 from 10.255.19.251
set protocols mpls label-switched-path pe2-pe3-p2mp-1 to 10.255.19.203
set protocols mpls label-switched-path pe2-pe3-p2mp-1 link-protection
set protocols mpls label-switched-path pe2-pe3-p2mp-1 p2mp vpn1-p2mp
set protocols mpls interface ae1.0
set protocols bgp group mpbg type internal
set protocols bgp group mpbg local-address 10.255.19.251
set protocols bgp group mpbg family inet unicast
set protocols bgp group mpbg family inet-vpn unicast per-prefix-label
set protocols bgp group mpbg family inet6 unicast
set protocols bgp group mpbg family inet6-vpn unicast per-prefix-label
set protocols bgp group mpbg family inet-mvpn signaling
set protocols bgp group mpbg family inet6-mvpn signaling
set protocols bgp group mpbg neighbor 10.255.19.253
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ldp interface ae1.0
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-4/2/5.0
set routing-instances vpn1 route-distinguisher 1:1
set routing-instances vpn1 provider-tunnel rsvp-te static-lsp vpn1-p2mp
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath
set routing-instances vpn1 protocols bgp group grp1 type external
set routing-instances vpn1 protocols bgp group grp1 export direct
set routing-instances vpn1 protocols bgp group grp1 peer-as 200
set routing-instances vpn1 protocols bgp group grp1 neighbor 198.51.100.4 family inet
    unicast
set routing-instances vpn1 protocols bgp group grp1 neighbor abcd:a:a:a:2::2 family inet6
    unicast
set routing-instances vpn1 protocols mvpn
set routing-options nonstop-routing
set routing-options autonomous-system 10
set routing-options forwarding-table export load-balancing-policy

```

```

CE2 set interfaces ge-0/0/5 unit 0 family inet address 198.51.100.4/24
    set interfaces ge-0/0/5 unit 0 family inet6 address abcd:a:a:a:2::2/126
    set protocols bgp group vpn1 type external
    set protocols bgp group vpn1 export direct
    set protocols bgp group vpn1 export vpn1
    set protocols bgp group vpn1 peer-as 10
    set protocols bgp group vpn1 neighbor 198.51.100.3 family inet unicast

```

```

set protocols bgp group vpn1 neighbor abcd:a:a:2::1 family inet6 unicast
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-options autonomous-system 200
set routing-options forwarding-table export load-balancing-policy

```

### Configuring Device PE1

#### 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*.

To configure Device PE1:

1. Specify the number of aggregated Ethernet interfaces to be created, configure the FPCs as vpn-core-facing-only, and enable enhanced IP network services.

```

[edit chassis]
user@PE1# set redundancy graceful-switchover
user@PE1# set aggregated-devices ethernet device-count 16
user@PE1# set fpc 8 vpn-localization vpn-core-facing-only
user@PE1# set network-services enhanced-ip

```

2. Configure the interfaces.

```

[edit interfaces]
user@PE1# set ge-2/0/0 unit 0 family inet address 192.0.2.1/24
user@PE1# set ge-2/0/0 unit 0 family inet6 address abcd:a:a:1::1/126

user@PE1# set ge-8/1/0 gigether-options 802.3ad ae0
user@PE1# set ge-8/1/9 gigether-options 802.3ad ae0

user@PE1# set ae0 unit 0 family inet address 192.0.2.3/24
user@PE1# set ae0 unit 0 family iso
user@PE1# set ae0 unit 0 family mpls

user@PE1# set lo0 unit 1 family inet address 203.0.113.1/24
user@PE1# set lo0 unit 1 family inet6 address abcd::10:0:1:1/128

```

3. Configure policy options to load balance the packets.

```

[edit policy-options policy-statement]
user@PE1# set direct from protocol direct
user@PE1# set direct then accept
user@PE1# set load-balancing-policy then load-balance per-packet

```

4. Configure the RSVP protocol on the interface.

```

[edit protocols rsvp]
user@PE1# set interface ae0.0

```

5. Configure the MPLS protocol.

```
[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set icmp-tunneling
user@PE1# set label-switched-path pe1-pe2-p2mp-1 from 10.255.19.254
user@PE1# set label-switched-path pe1-pe2-p2mp-1 to 10.255.19.251
user@PE1# set label-switched-path pe1-pe2-p2mp-1 link-protection
user@PE1# set label-switched-path pe1-pe2-p2mp-1 p2mp vpn1-p2mp
user@PE1# set label-switched-path pe1-pe3-p2mp-1 from 10.255.19.254
user@PE1# set label-switched-path pe1-pe3-p2mp-1 to 10.255.19.203
user@PE1# set label-switched-path pe1-pe3-p2mp-1 link-protection
user@PE1# set label-switched-path pe1-pe3-p2mp-1 p2mp vpn1-p2mp
user@PE1# set interface ae0.0
```

6. Configure the BGP protocol for the mpbg group.

```
[edit protocols bgp group mpbg]
user@PE1# set type internal
user@PE1# set local-address 10.255.19.254
user@PE1# set family inet unicast
user@PE1# set family inet-vpn unicast
user@PE1# set family inet6 unicast
user@PE1# set family inet6-vpn unicast
user@PE1# set family inet-mvpn signaling
user@PE1# set family inet6-mvpn signaling
user@PE1# set neighbor 10.255.19.253
```

7. Configure the OSPF protocol.

```
[edit protocols ospf]
user@PE1# set traffic-engineering
user@PE1# set area 0.0.0.0 interface ae0.0
user@PE1# set area 0.0.0.0 interface lo0.0 passive
```

8. Configure the LDP protocol on the interface.

```
[edit protocols]
user@PE1# set ldp interface ae0.0
```

9. Create an instance type and configure the routing instances on the interface.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-2/0/0.0
user@PE1# set interface lo0.1
```

10. Configure the route distinguisher, and configure the static LSP for the provider tunnel RSVP-TE.

```
[edit routing-instances vpn1]
user@PE1# set route-distinguisher 1:1
user@PE1# set provider-tunnel rsvp-te static-lsp vpn1-p2mp
```

11. Configure the VRF target and the VRF target label for the routing instance.

```
[edit routing-instances vpn1]
user@PE1# set vrf-target target:1:1
user@PE1# set vrf-table-label
```

12. Configure the multipath routing option for a routing instance, and configure the localized fib routing option for the routing instance.

```
[edit routing-instances vpn1 routing-options]
user@PE1# set multipath
user@PE1# set localized-fib
```

13. Configure the group of BGP protocols for a routing instance.

```
[edit routing-instances vpn1 protocols bgp group grp1]
user@PE1# set type external
user@PE1# set export direct
user@PE1# set peer-as 100
user@PE1# set neighbor 192.0.2.2 family inet unicast
user@PE1# set neighbor abcd:a:a:1::2 family inet6 unicast
```

14. Configure the MVPN protocols.

```
[edit routing-instances vpn1]
user@PE1# set protocols mvpn
```

15. Configure the nonstop active routing and the autonomous system number for a routing option.

```
[edit routing-options]
user@PE1# set nonstop-routing
user@PE1# set autonomous-system 10
```

16. Configure the load-balancing policy for the forwarding table and extended space for the chained composite next hop for the L3VPN of the forwarding table.

```
[edit routing-options]
user@PE1# set forwarding-table export load-balancing-policy
user@PE1# set forwarding-table chained-composite-next-hop ingress l3vpn
extended-space
```

---

## Results

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

```
user@PE1# show chassis
redundancy {
  graceful-switchover;
}
aggregated-devices {
```

```

    ethernet {
        device-count 16;
    }
}
fpc 8 {
    vpn-localization vpn-core-facing-only;
}
network-services enhanced-ip;

user@PE1# show interfaces
ge-2/0/0 {
    unit 0 {
        family inet {
            address 192.0.2.1/24;
        }
        family inet6 {
            address abcd:a:a:a::1/126;
        }
    }
}
ge-8/1/0 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-8/1/9 {
    gigether-options {
        802.3ad ae0;
    }
}
ae0 {
    unit 0 {
        family inet {
            address 192.0.2.3/24;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 1 {
        family inet {
            address 203.0.113.1/24;
        }
        family inet6 {
            address abcd::10:0:1:1/128;
        }
    }
}

user@PE1# show policy-options
policy-statement direct {
    from protocol direct;
    then accept;
}
policy-statement load-balancing-policy {
    then {

```

```
        load-balance per-packet;
    }
}

user@PE1# show routing-options
nonstop-routing;
autonomous-system 10;
forwarding-table {
    export load-balancing-policy;
    chained-composite-next-hop {
        ingress {
            l3vpn extended-space;
        }
    }
}

user@PE1# show routing-instances
vpn1 {
    instance-type vrf;
    interface ge-2/0/0.0;
    interface lo0.1;
    route-distinguisher 1:1;
    provider-tunnel {
        rsvp-te {
            static-lsp vpn1-p2mp;
        }
    }
    vrf-target target:1:1;
    vrf-table-label;
    routing-options {
        multipath;
        localized-fib;
    }
    protocols {
        bgp {
            group grp1 {
                type external;
                export direct;
                peer-as 100;
                neighbor 192.0.2.2 {
                    family inet {
                        unicast;
                    }
                }
                neighbor abcd:a:a:a:1::2 {
                    family inet6 {
                        unicast;
                    }
                }
            }
        }
    }
    mvpn;
}

user@PE1# show protocols
rsvp {
```



```

    interface ae0.0;
  }
  mpls {
    ipv6-tunneling;
    icmp-tunneling;
    label-switched-path pe1-pe2-p2mp-1 {
      from 10.255.19.254;
      to 10.255.19.251;
      link-protection;
      p2mp vpn1-p2mp;
    }
    label-switched-path pe1-pe3-p2mp-1 {
      from 10.255.19.254;
      to 10.255.19.203;
      link-protection;
      p2mp vpn1-p2mp;
    }
  }
  interface ae0.0;
}
bgp {
  group mppbg {
    type internal;
    local-address 10.255.19.254;
    family inet {
      unicast;
    }
    family inet-vpn {
      unicast;
    }
    family inet6 {
      unicast;
    }
    family inet6-vpn {
      unicast;
    }
    family inet-mvpn {
      signaling;
    }
    family inet6-mvpn {
      signaling;
    }
    neighbor 10.255.19.253;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ae0.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ae0.0;
}

```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying VRF Localization on page 158](#)
- [Verifying VRF Localization for a VPN on page 158](#)

---

### Verifying VRF Localization

**Purpose** Verify the localization of VRF in a Layer 3 VPN.

**Action** From operational mode, run the **show route vpn-localization** command for Device PE1.

```
user@PE1> show route vpn-localization
```

```
Routing table: vpn1.inet, Localized
  Index: 7, Address Family: inet, Localization status: Complete
  Local FPC's: 2 8
```

```
Routing table: vpn1.inet6, Localized
  Index: 7, Address Family: inet6, Localization status: Complete
  Local FPC's: 2 8
```

```
Routing table: vpn2.inet, Non-localized
  Index: 8, Address Family: inet, Localization status: Complete
  Local FPC's: All
```

```
Routing table: vpn2.inet6, Non-localized
  Index: 8, Address Family: inet6, Localization status: Complete
  Local FPC's: All
```

**Meaning** The output shows the localization information of all the VRFs.

---

### Verifying VRF Localization for a VPN

**Purpose** Verify VRF localization for a VPN.

**Action** From operational mode, run the **show route vpn-localization vpn-name *vpn-name*** command.

```
user@PE1> show route vpn-localization vpn-name vpn1
```

```
Routing table: vpn1.inet, Localized
  Index: 7, Address Family: inet, Localization status: Complete
  Local FPC's: 2 8
```

```
Routing table: vpn1.inet6, Localized
  Index: 7, Address Family: inet6, Localization status: Complete
  Local FPC's: 2 8
```

**Meaning** The output shows the VPN localization of a VPN.

- Related Documentation**
- [Understanding VRF Localization in Layer 3 VPNs on page 145](#)
  - [localized-fib on page 1095](#)
  - [vpn-localization on page 1150](#)

## Filtering Packets in Layer 3 VPNs Based on IP Headers

Including the **vrf-table-label** statement in the configuration for a routing instance makes it possible to map the inner label to a specific VRF routing table; such mapping allows the examination of the encapsulated IP header at an egress VPN router. You might want to enable this functionality so that you can do either of the following:

- Forward traffic on a PE-router-to-CE-device interface, in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch).

The first lookup is done on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared medium.

- Perform egress filtering at the egress PE router.

The first lookup on the VPN label is done to determine which VRF routing table to refer to, and the second lookup is done on the IP header to determine how to filter and forward packets. You can enable this functionality by configuring output filters on the VRF interfaces.

When you include the **vrf-table-label** statement in the configuration of a VRF routing table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table. Any routes in such a VRF routing table are advertised with the LSI logical interface label allocated for the VRF routing table. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table.

To filter traffic based on the IP header, include the **vrf-table-label** statement:

```
vrf-table-label {  
    source-class-usage;  
}
```

You can include the statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

You can include the **vrf-table-label** statement for both IPv4 and IPv6 Layer 3 VPNs. If you include the statement for a dual-stack VRF routing table (where both IPv4 and IPv6 routes are supported), the statement applies to both the IPv4 and IPv6 routes and the same label is advertised for both sets of routes.

You can also configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement by also including the **source-class-usage** option. Include the **source-class-usage** statement at the **[edit routing-instances *routing-instance-name* vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the **vrf** instance type (Layer 3 VPNs). DCU is not supported for the **vrf-table-label** statement. For more information, see *Enabling Source Class and Destination Class Usage*.

The following sections provide more information about traffic filtering based on the IP header:

- [Egress Filtering Options on page 160](#)
- [Support on Aggregated and VLAN Interfaces for IP-Based Filtering on page 161](#)
- [Support on ATM and Frame Relay Interfaces for IP-Based Filtering on page 161](#)
- [Support on Ethernet, SONET/SDH, and T1/T3/E3 Interfaces for IP-Based Filtering on page 162](#)
- [Support on SONET/SDH and DS3/E3 Channelized Enhanced Intelligent Queuing Interfaces for IP-Based Filtering on page 162](#)
- [Support on Multilink PPP and Multilink Frame Relay Interfaces for IP-Based Filtering on page 164](#)
- [Support for IP-Based Filtering of Packets with Null Top Labels on page 164](#)
- [General Limitations on IP-Based Filtering on page 165](#)

## Egress Filtering Options

You can enable egress filtering (which allows egress Layer 3 VPN PE routers to perform lookups on the VPN label and IP header at the same time) by including the **vrf-table-label** statement at the **[edit routing-instances *instance-name*]** hierarchy level. There is no restriction on including this statement for CE-router-to-PE-router interfaces, but there are several limitations on other interface types, as described in subsequent sections in this topic.

You can also enable egress filtering by configuring a VPN tunnel (VT) interface on routing platforms equipped with a Tunnel Services Physical Interface Card (PIC). When you enable egress filtering this way, there is no restriction on the type of core-facing interface used. There is also no restriction on the type of CE-router-to-PE-router interface used.

## Support on Aggregated and VLAN Interfaces for IP-Based Filtering

Support for the **vrf-table-label** statement over aggregated and VLAN interfaces is available on the routers summarized in [Table 4 on page 161](#).

*Table 4: Support for Aggregated and VLAN Interfaces*

Interfaces	M Series Router Without an Enhanced FPC	M Series Router with an Enhanced FPC	M320 Router	T Series Router
Aggregated	No	Yes	Yes	Yes
VLAN	No	Yes	Yes	Yes



**NOTE:** The **vrf-table-label** statement is not supported for Aggregated Gigabit Ethernet, 10-Gigabit Ethernet, and VLAN physical interfaces on M120 routers.

## Support on ATM and Frame Relay Interfaces for IP-Based Filtering

Support for the **vrf-table-label** statement over Asynchronous Transfer Mode (ATM) and Frame Relay interfaces is available on the routers summarized in [Table 5 on page 161](#).

*Table 5: Support for ATM and Frame Relay Interfaces*

Interfaces	M Series Router Without an Enhanced FPC	M Series Router with an Enhanced FPC	M320 Router	T Series Router
ATM1	No	No	No	No
ATM2 intelligent queuing (IQ)	No	Yes	Yes	Yes
Frame Relay	No	Yes	Yes	Yes
Channelized	No	No	No	No

When you include the **vrf-table-label** statement, be aware of the following limitations with ATM or Frame Relay interfaces:

- The **vrf-table-label** statement is supported on ATM interfaces, but with the following limitations:
  - ATM interfaces can be configured on the M320 router and the T Series routers, and on M Series routers with an enhanced FPC.
  - The interface can only be a PE router interface receiving traffic from a P router.
  - The router must have an ATM2 IQ PIC.
- The **vrf-table-label** statement is also supported on Frame Relay encapsulated interfaces, but with the following limitations:
  - Frame Relay interfaces can be configured on the M320 router and the T Series routers, and on M Series routers with an enhanced FPC.
  - The interface can only be a PE router interface receiving traffic from a P router.

### Support on Ethernet, SONET/SDH, and T1/T3/E3 Interfaces for IP-Based Filtering

Support for the **vrf-table-label** statement over Ethernet, SONET/SDH, and T1/T3/E3 interfaces is available on the routers summarized in [Table 6 on page 162](#).

*Table 6: Support for Ethernet, SONET/SDH, and T1/T3/E3 Interfaces*

Interfaces	M Series Router Without an Enhanced FPC	M Series Router with an Enhanced FPC	M320 Router	T Series Router
Ethernet	Yes	Yes	Yes	Yes
SONET/SDH	Yes	Yes	Yes	Yes
T1/T3/E3	Yes	Yes	Yes	Yes

Only the following Ethernet PICs support the **vrf-table-label** statement on M Series routers without an Enhanced FPC:

- 1-port Gigabit Ethernet
- 2-port Gigabit Ethernet
- 4-port Fast Ethernet

### Support on SONET/SDH and DS3/E3 Channelized Enhanced Intelligent Queuing Interfaces for IP-Based Filtering

Support for the **vrf-table-label** statement for the specified channelized IQE interfaces is only available on M120 and M320 routers with Enhanced III FPCs as summarized in [Table 7 on page 163](#).

*Table 7: Support for Channelized IQE Interfaces on M320 Routers with Enhanced III FPCs*

Interfaces	M120 Routers with Enhanced III FPCs	M320 Routers with Enhanced III FPCs
OC12	Yes	Yes
STM4	Yes	Yes
OC3	Yes	Yes
STM1	Yes	Yes
DS3	Yes	Yes
E3	Yes	Yes

The following IQE Type-1 PICs are supported:

- 1-port OC12/STM4 IQE with SFP
- 4-port OC3/STM1 IQE with SFP
- 4-port DS3/E3 IQE with BNC
- 2-port Channelized OC3/STM1 IQE with SFP, with no SONET partitions
- 1-port Channelized OC12/STM4 IQE with SFP, with no SONET partitions

The following constraints are applicable with respect to a router configuration utilizing logical systems:

- Multipoint IQE PIC interfaces constraints—On multipoint IQE PICs, such as the 2-port Channelized OC3/STM1 IQE with SFP, if the port 1 interface is configured as one logical system with its own routing-instance and the port 2 interface is configured as a different logical system with its own routing instances such that there are core-facing logical interfaces on both port 1 and port 2, then you cannot configure the **vrf-table-label** statement on routing-instance in both logical systems. Only one set of LSI labels are supported; the last routing instance with the **vrf-table-label** statement configured is committed.
- Frame Relay encapsulation and logical interfaces across logical systems constraints—Similar to the multipoint PIC with logical systems, if you try to configure one logical interface of an IQE PIC with Frame Relay encapsulation in one logical system and configure another logical interface on the same IQE PIC in the second logical system, the configuration will not work for all the **vrf-table-label** statement configured instances. It will only work for the instances configured in one of the logical systems.

Both the above constraints occur because the router configuration maintains one LSI tree in the Packet Forwarding Engine per logical system, which is common across all streams. The stream channel table lookup is then adjusted to point to the LSI tree. In the case of multipoint type-1 IQE PICs, all physical interfaces share the same stream. Therefore, the logical interfaces (multipoint or not) obviously share the same stream. Consequently,

the LSI binding is at the stream level. Hence, provisioning logical interfaces under the same stream provisioned to be core-facing and supporting a different set of routing instances with the **vrf-table-label** statement is not supported.

## Support on Multilink PPP and Multilink Frame Relay Interfaces for IP-Based Filtering

Support for the **vrf-table-label** statement over Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR) interfaces is available on the routers summarized in [Table 8 on page 164](#).

*Table 8: Support for Multilink PPP and Multilink Frame Relay Interfaces*

Interfaces	M Series Router Without an Enhanced FPC	M Series Router with an Enhanced FPC	M320	T Series Router	MX Series Router
MLPPP	No	Yes	No	No	No
End-to-End MLFR (FRF.15)	No	Yes	No	No	No
UNI/NNI MLFR (FRF.16)	No	No	No	No	No

M Series routers must have an AS PIC to support the **vrf-table-label** statement over MLPPP and MLFR interfaces. The **vrf-table-label** statement over MLPPP interfaces is not supported on M120 routers.

## Support for IP-Based Filtering of Packets with Null Top Labels

You can include the **vrf-table-label** statement in the configuration for core-facing interfaces receiving MPLS packets with a null top label, which might be transmitted by some vendors' equipment. These packets can be received only on the M320 router, the M10i router, and T Series Core routers using one of the following PICs:

- 1-port Gigabit Ethernet with SFP
- 2-port Gigabit Ethernet with SFP
- 4-port Gigabit Ethernet with SFP
- 10-port Gigabit Ethernet with SFP
- 1-port SONET STM4
- 4-port SONET STM4
- 1-port SONET STM16
- 1-port SONET STM16 (non-SFP)
- 4-port SONET STM16
- 1-port SONET STM64



The following PICs can receive packets with null top labels, but only when installed in an M120 router or an M320 router with an Enhanced III FPC:

- 1-port 10-Gigabit Ethernet
- 1-port 10-Gigabit Ethernet IQ2

## General Limitations on IP-Based Filtering

The following limitations apply when you include the **vrf-table-label** statement:

- Firewall filters cannot be applied to interfaces included in a routing instance on which you have configured the **vrf-table-label** statement.
- The time-to-live (TTL) value in the MPLS header is not copied back to the IP header of packets sent from the PE router to the CE router.
- You cannot include the **vrf-table-label** statement in a routing instance configuration that also includes a virtual loopback tunnel interface; the commit operation fails in this case.
- When you include the statement, MPLS packets with label-switched interface (LSI) labels that arrive on core-facing interfaces are not counted at the logical interface level if the core-facing interface is any of the following:
  - ATM
  - Frame Relay
  - Ethernet configured with VLANs
  - Aggregated Ethernet configured with VLANs
- For LMNR, Stoli, and I-Chip-based Packet Forwarding Engines, you cannot include the statement in the configuration of a VRF routing instance if the PE-router-to-P-router interface is any of the following interfaces:



**NOTE:** The **vrf-table-label** statement is supported when the PE-router-to-P-router interface is a tunnel interface on a Junos Trio-based Packet Forwarding Engine, so no limitation applies.

- Aggregated SONET/SDH interface
- Channelized interface
- Tunnel interface (for example, generic routing encapsulation [GRE] or IP Security [IPsec])
- Circuit cross-connect (CCC) or translational cross-connect (TCC) encapsulated interface
- Logical tunnel interface
- Virtual private LAN service (VPLS) encapsulated interface



**NOTE:** All CE-router-to-PE-router and PE-router-to-CE-router interfaces are supported.

- You cannot include the **vrf-table-label** statement in the configuration of a VRF routing instance if the PE-router-to-P-router PIC is one of the following PICs:
  - 10-port E1
  - 8-port Fast Ethernet
  - 12-port Fast Ethernet
  - 48-port Fast Ethernet
  - ATM PIC other than the ATM2 IQ
- Label-switched interface (LSI) traffic statistics are not supported for Intelligent Queuing 2 (IQ2), Enhanced IQ (IQE), and Enhanced IQ2 (IQ2E) PICs on M Series routers.

**Related Documentation**

- *Enabling Source Class and Destination Class Usage*

## Configuring a Label Allocation and Substitution Policy for VPNs

You can control label-advertisements on MPLS ingress and AS border routers (ASBRs). Labels can be assigned on a per-next-hop (by default) or on a per-table basis (by configuring the **vrf-table-label** statement). This choice affects all routes of a given routing instance. You can also configure a policy to generate labels on a per-route basis by specifying a label allocation policy.

To specify a label allocation policy for the routing instance, configure the **label** statement and specify a label allocation policy using the **allocation** option:

```
label {
    allocation label-allocation-policy;
}
```

You can configure this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* routing-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

To configure the label allocation policy, include the **label-allocation** statement at the [edit policy-options policy-statement *policy-statement-name* term *term-name* then]

hierarchy level. You can configure the label allocation mode as either **per-nexthop** or **per-table**.

For a VPN option B ASBR, labels for transit routes are substituted for a local virtual tunnel label or vrf-table-label label. When a VRF table is configured on the ASBR (this type of configuration is uncommon for the option B model), the ASBR does not generate MPLS swap or swap and push state for transit routes. Instead, the ASBR re-advertises a local virtual-tunnel or vrf-table-label label and forwards that transit traffic based on IP forwarding tables. The label substitution helps to conserve labels on Juniper Networks routers.

However, this type of label substitution effectively breaks the MPLS forwarding path, which becomes visible when using an MPLS OAM command such as LSP ping. You can configure the way in which labels are substituted on a per-route basis by specifying a label substitution policy.

To specify a label substitution policy for the routing instance, configure the **label** statement and specify a label substitution policy using the **substitution** option:

```
label {
  substitution label-substitution-policy;
}
```

You can configure this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* routing-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

The label substitution policy is used to determine whether or not a label should be substituted on an ASBR router. The results of the policy operation are either **accept** (label substitution is performed) or **reject** (label substitution is not performed). The default behavior is **accept**. The following set command example illustrates how you can configure a **reject** label substitution policy: **set policy-options policy-statement no-label-substitution term default then reject**.



# Distributing Routes in Layer 3 VPNs

- [Configuring Routing Between PE and CE Routers in Layer 3 VPNs on page 169](#)
- [Configuring an OSPF Domain ID for a Layer 3 VPN on page 179](#)
- [Example: Giving VPN Backbones Priority with OSPFv2 Sham Links for Layer 3 VPNs on page 185](#)
- [Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs on page 195](#)
- [Configuring an LDP-over-RSVP VPN Topology on page 196](#)

## Configuring Routing Between PE and CE Routers in Layer 3 VPNs

---

For the PE router to distribute VPN-related routes to and from connected CE routers, you must configure routing within the VPN routing instance. You can configure a routing protocol—BGP, OSPF, or RIP—or you can configure static routing. For the connection to each CE router, you can configure only one type of routing.

The following sections explain how to configure VPN routing between the PE and CE routers:

- [Configuring BGP Between the PE and CE Routers on page 169](#)
- [Configuring OSPF Between the PE and CE Routers on page 170](#)
- [Configuring OSPF Sham Links for Layer 3 VPNs on page 171](#)
- [Configuring an OSPF Domain ID on page 174](#)
- [Configuring RIP Between the PE and CE Routers on page 177](#)
- [Configuring Static Routes Between the PE and CE Routers on page 179](#)

## Configuring BGP Between the PE and CE Routers

To configure BGP as the routing protocol between the PE and the CE routers, include the **bgp** statement:

```
bgp {  
  group group-name {  
    peer-as as-number;  
    neighbor ip-address;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

Please be aware of the following limitations regarding configuring BGP for routing instances:

- In a VRF routing instance, do not configure the local autonomous system (AS) number using an AS number that is already in use by a remote BGP peer in a separate VRF routing instance. Doing so creates an autonomous system loop where all the routes received from this remote BGP peer are hidden.

You configure the local AS number using either the **autonomous-system** statement at the [edit routing-instances *routing-instance-name* routing-options] hierarchy level or the **local-as** statement at any of the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols bgp]
- [edit routing-instances *routing-instance-name* protocols bgp group *group-name*]
- [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]

You configure the AS number for a BGP peer using the **peer-as** statement at the [edit routing-instances *routing-instance-name* protocols bgp group *group-name*] hierarchy level.

## Configuring OSPF Between the PE and CE Routers

You can configure OSPF (version 2 or version 3) to distribute VPN-related routes between PE and CE routers.

The following sections describe how to configure OSPF as a routing protocol between the PE and the CE routers:

- [Configuring OSPF Version 2 Between the PE and CE Routers on page 170](#)
- [Configuring OSPF Version 3 Between the PE and CE Routers on page 171](#)

---

### Configuring OSPF Version 2 Between the PE and CE Routers

To configure OSPF version 2 as the routing protocol between a PE and CE router, include the **ospf** statement:

```
ospf {  
  area area {  
    interface interface-name;  
  }  
}
```

```
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

### Configuring OSPF Version 3 Between the PE and CE Routers

To configure OSPF version 3 as the routing protocol between a PE and CE router, include the **ospf3** statement:

```
ospf3 {
  area area {
    interface interface-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

### Configuring OSPF Sham Links for Layer 3 VPNs

When you configure OSPF between the PE and CE routers of a Layer 3 VPN, you can also configure OSPF sham links to compensate for issues related to OSPF intra-area links.

The following sections describe OSPF sham links and how to configure them:

- [OSPF Sham Links Overview on page 171](#)
- [Configuring OSPF Sham Links on page 172](#)
- [OSPF Sham Links Example on page 173](#)


#### OSPF Sham Links Overview

[Figure 15 on page 172](#) provides an illustration of when you might configure an OSPF sham link. Router CE1 and Router CE2 are located in the same OSPF area. These CE routers are linked together by a Layer 3 VPN over Router PE1 and Router PE2. In addition, Router CE1 and Router CE2 are connected by an intra-area link used as a backup.

OSPF treats the link through the Layer 3 VPN as an interarea link. By default, OSPF prefers intra-area links to interarea links, so OSPF selects the backup intra-area link as the active path. This is not acceptable in configurations where the intra-area link is not the expected primary path for traffic between the CE routers.

An OSPF sham link is also an intra-area link, except that it is configured between the PE routers as shown in [Figure 15 on page 172](#). You can configure the metric for the sham link to ensure that the path over the Layer 3 VPN is preferred to a backup path over an intra-area link connecting the CE routers.

**Figure 15: OSPF Sham Link**



You should configure an OSPF sham link under the following circumstances:

- Two CE routers are linked together by a Layer 3 VPN.
- These CE routers are in the same OSPF area.
- An intra-area link is configured between the two CE routers.

If there is no intra-area link between the CE routers, you do not need to configure an OSPF sham link.

For more information about OSPF sham links, see the Internet draft [draft-ietf-l3vpn-ospf-2547-01.txt](#), *OSPF as the PE/CE Protocol in BGP/MPLS VPNs*.

### Configuring OSPF Sham Links

The sham link is an unnumbered point-to-point intra-area link and is advertised by means of a type 1 link-state advertisement (LSA). Sham links are valid only for routing instances and OSPF version 2.

Each sham link is identified by a combination of the local and remote sham link end-point address and the OSPF area to which it belongs. Sham links must be configured manually. You configure the sham link between two PE routers, both of which are within the same VRF routing instance.

You need to specify the address for the local end point of the sham link. This address is used as the source for the sham link packets and is also used by the remote PE router as the sham link remote end-point.

The OSPF sham link's local address must be specified with a loopback address for the local VPN. The route to this address must be propagated by BGP. Specify the address for the local end point using the **local** option of the **sham-link** statement:



```
sham-link {
  local address;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols ospf]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf]

The OSPF sham link's remote address must be specified with a loopback address for the remote VPN. The route to this address must be propagated by BGP. To specify the address for the remote end point, include the **sham-link-remote** statement:

```
sham-link-remote address <metric number>;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols ospf area *area-id*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id*]

Optionally, you can include the **metric** option to set a metric value for the remote end point. The metric value specifies the cost of using the link. Routes with lower total path metrics are preferred over those with higher path metrics.

You can configure a value from 1 through 65,535. The default value is 1.

### OSPF Sham Links Example

This example shows how to enable OSPF sham links on a PE router.

The following is the loopback interface configuration on the PE router. The address configured is for the local end point of the OSPF sham link:

```
[edit]
interfaces {
  lo0 {
    unit 1 {
      family inet {
        address 10.1.1.1/32;
      }
    }
  }
}
```

The following is the routing instance configuration on the PE router, including the configuration for the OSPF sham link. The **sham-link local** statement is configured with the address for the local loopback interface:

```
[edit]
routing-instances {
  example-sham-links {
    instance-type vrf;
    interface e1-1/0/2.0;
```

```

interface lo0.1;
route-distinguisher 3:4;
vrf-import vpn-red-import;
vrf-export vpn-red-export;
protocols {
  ospf {
    sham-link local 10.1.1.1;
    area 0.0.0.0 {
      sham-link-remote 10.2.2.2 metric 1;
      interface e1-1/0/2.0 metric 1;
    }
  }
}
}

```

## Configuring an OSPF Domain ID

For most OSPF configurations involving Layer 3 VPNs, you do not need to configure an OSPF domain ID. However, for a Layer 3 VPN connecting multiple OSPF domains, configuring OSPF domain IDs can help you control LSA translation (for Type 3 and Type 5 LSAs) between the OSPF domains and back-door paths. Each VPN routing and forwarding (VRF) table in a PE router associated with an OSPF instance is configured with the same OSPF domain ID. The default OSPF domain ID is the null value 0.0.0.0. As shown in [Table 9 on page 174](#), a route with a null domain ID is handled differently from a route without any domain ID at all.

**Table 9: How a PE Router Redistributes and Advertises Routes**

Route Received	Domain ID of the Route Received	Domain ID on the Receiving Router	Route Redistributed and Advertised As
Type 3 route	A.B.C.D	A.B.C.D	Type 3 LSA
Type 3 route	A.B.C.D	E.F.G.H	Type 5 LSA
Type 3 route	0.0.0.0	0.0.0.0	Type 3 LSA
Type 3 route	Null	0.0.0.0	Type 3 LSA
Type 3 route	Null	Null	Type 3 LSA
Type 3 route	0.0.0.0	Null	Type 3 LSA
Type 3 route	A.B.C.D	Null	Type 5 LSA
Type 3 route	Null	A.B.C.D	Type 5 LSA
Type 5 route	Not applicable	Not applicable	Type 5 LSA

You can configure an OSPF domain ID for both version 2 and version 3 of OSPF. The only difference in the configuration is that you include statements at the `[edit routing-instances`

***routing-instance-name protocols ospf***] hierarchy level for OSPF version 2 and at the **[edit routing-instances *routing-instance-name* protocols ospf3]** hierarchy level for OSPF version 3. The configuration descriptions that follow present the OSPF version 2 statement only. However, the substatements are also valid for OSPF version 3.

To configure an OSPF domain ID, include the **domain-id** statement:

```
domain-id domain-id;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols ospf]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf]**

You can set a VPN tag for the OSPF external routes generated by the PE router to prevent looping. By default, this tag is automatically calculated and needs no configuration. However, you can configure the domain VPN tag for Type 5 LSAs explicitly by including the **domain-vpn-tag** statement:

```
no-domain-vpn-tag number;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols ospf]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf]**

The range is 1 through 4,294,967,295 ( $2^{32} - 1$ ). If you set VPN tags manually, you must set the same value for all PE routers in the VPN.

For an example of this type of configuration, see [“Configuring an OSPF Domain ID for a Layer 3 VPN” on page 179](#).

### Hub-and-Spoke Layer 3 VPNs and OSPF Domain IDs

---

The default behavior of an OSPF domain ID causes some problems for hub-and-spoke Layer 3 VPNs configured with OSPF between the hub PE router and the hub CE router when the routes are not aggregated. A hub-and-spoke configuration has a hub PE router with direct links to a hub CE router. The hub PE router receives Layer 3 BGP updates from the other remote spoke PE routers, and these are imported into the spoke routing instance. From the spoke routing instance, the OSPF LSAs are originated and sent to the hub CE router.

The hub CE router typically aggregates these routes, and then sends these newly originated LSAs back to the hub PE router. The hub PE router exports the BGP updates to the remote spoke PE routers containing the aggregated prefixes. However, if there are nonaggregated Type 3 summary LSAs or external LSAs, two issues arise with regard to how the hub PE router originates and sends LSAs to the hub CE router, and how the hub PE router processes LSAs received from the hub CE router:

- By default, all LSAs originated by the hub PE router in the spoke routing instance have the DN bit set. Also, all externally originated LSAs have the VPN route tag set. These settings help prevent routing loops. For Type 3 summary LSAs, routing loops are not a concern because the hub CE router, as an area border router (ABR), reoriginates the LSAs with the DN bit clear and sends them back to the hub PE router. However, the hub CE router does not reoriginate external LSAs, because they have an AS flooding scope.

You can originate the external LSAs (before sending them to the hub CE router) with the DN bit clear and the VPN route tag set to 0 by altering the hub PE router's routing instance configuration. To clear the DN bit and set the VPN route tag to zero on external LSAs originated by a PE router, configure 0 for the **domain-vpn-tag** statement at the **[edit routing-instances routing-instance-name protocols ospf]** hierarchy level. You should include this configuration in the routing instance on the hub PE router facing the hub CE router where the LSAs are sent. When the hub CE router receives external LSAs from the hub PE router and then forwards them back to the hub PE router, the hub PE router can use the LSAs in its OSPF route calculation.

- When LSAs flooded by the hub CE router arrive at the hub PE router's routing instance, the hub PE router, acting as an ABR, does not consider these LSAs in its OSPF route calculations, even though the LSAs do not have the DN bits set and the external LSAs do not have a VPN route tag set. The LSAs are assumed to be from a disjoint backbone area.

You can change the configuration of the PE router's routing instance to cause the PE router to act as a non-ABR by including the **disable** statement at the **[edit routing-instances routing-instance-name protocols ospf domain-id]** hierarchy level. You make this configuration change to the hub PE router that receives the LSAs from the hub CE router.

By making this configuration change, the PE router's routing instance acts as a non-ABR. The PE router then considers the LSAs arriving from the hub CE router as if they were coming from a contiguous nonbackbone area.

## Configuring RIP Between the PE and CE Routers

For a Layer 3 VPN, you can configure RIP on the PE router to learn the routes of the CE router or to propagate the routes of the PE router to the CE router. RIP routes learned from neighbors configured at any **[edit routing-instances]** hierarchy level are added to the routing instance's **inet** table (*instance\_name.inet.0*).

To configure RIP as the routing protocol between the PE and the CE router, include the **rip** statement:

```
rip {
  group group-name {
    export policy-names;
    neighbor interface-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances routing-instance-name protocols]**
- **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols]**



**NOTE:** The **[edit logical-systems]** hierarchy level is not applicable in ACX Series routers.

By default, RIP does not advertise the routes it receives. To advertise routes from a PE router to a CE router, you need to configure an export policy on the PE router for RIP. For information about how to define policies for RIP, see *Example: Applying Policies to RIP Routes Imported from Neighbors*.

To specify an export policy for RIP, include the **export** statement:

```
export [ policy-names ];
```

You can include this statement for RIP at the following hierarchy levels:

- **[edit 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]**



**NOTE:** The **[edit logical-systems]** hierarchy level is not applicable in ACX Series routers.

To install routes learned from a RIP routing instance into multiple routing tables, include the **rib-group** and **group** statements:

```
rib-group inet group-name;
group group-name {
```

```
neighbor interface-name;  
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

To configure a routing table group, include the **rib-groups** statement:

```
rib-groups group-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

To add a routing table to a routing table group, include the **import-rib** statement. The first routing table name specified under the **import-rib** statement must be the name of the routing table you are configuring. For more information about how to configure routing tables and routing table groups, see *Junos OS Routing Protocols Library*.

```
import-rib [ group-names ];
```

You can include this statement at the following hierarchy levels:

- [edit routing-options rib-groups *group-name*]
- [edit logical-systems *logical-system-name* routing-options rib-groups *group-name*]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

RIP instances are supported only for VRF instance types. You can configure multiple instances of RIP for VPN support only. You can use RIP in the customer edge-provider edge (CE-PE) environment to learn routes from the CE router and to propagate the PE router's instance routes in the CE router.

RIP routes learned from neighbors configured under any instance hierarchy are added to the instance's routing table, *instance-name.inet.0*.

RIP does not support routing table groups; therefore, it cannot import routes into multiple tables as the OSPF or OSPFv3 protocol does.

## Configuring Static Routes Between the PE and CE Routers

You can configure static (nonchanging) routes between the PE and CE routers of a VPN routing instance. To configure a static route for a VPN, you need to configure it within the VPN routing instance configuration at the **[edit routing-instances routing-instance-name routing-options]** hierarchy level.

To configure a static route between the PE and the CE routers, include the **static** statement:

```
static {
  route destination-prefix {
    next-hop [ next-hops ];
    static-options;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances routing-instance-name routing-options]**
- **[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options]**



**NOTE:** The **[edit logical-systems]** hierarchy level is not applicable in ACX Series routers.

For more information about configuring routing protocols and static routes, see *Junos OS Routing Protocols Library*.

## Configuring an OSPF Domain ID for a Layer 3 VPN


This example illustrates how to configure an OSPF domain ID for a VPN by using OSPF as the routing protocol between the PE and CE routers. Routes from an OSPF domain need an OSPF domain ID when they are distributed in BGP as VPN-IPv4 routes in VPNs with multiple OSPF domains. In a VPN connecting multiple OSPF domains, the routes from one domain might overlap with the routes of another.

The domain ID that is configured in a routing instance identifies the OSPF domain and is used to identify the route origination. The domain ID that is configured on a community policy is used in setting exported routes.

For more information about OSPF domain IDs and Layer 3 VPNs, see [“Configuring an OSPF Domain ID” on page 174](#).

Figure 16 on page 180 shows this example's configuration topology. Only the configuration for Router PE1 is provided. The configuration for Router PE2 can be similar to the configuration for Router PE1. There are no special configuration requirements for the CE routers.

Figure 16: Example of a Configuration Using an OSPF Domain ID



For configuration information, see the following sections:

- [Configuring Interfaces on Router PE1 on page 180](#)
- [Configuring Routing Options on Router PE1 on page 181](#)
- [Configuring Protocols on Router PE1 on page 181](#)
- [Configuring Policy Options on Router PE1 on page 181](#)
- [Configuring the Routing Instance on Router PE1 on page 182](#)
- [Configuration Summary for Router PE1 on page 183](#)

## Configuring Interfaces on Router PE1

You need to configure two interfaces for Router PE1—the **so-0/0/0** interface for traffic to Router CE1 (San Francisco) and the **so-0/0/1** interface for traffic to a P router in the service provider's network.

Configure the interfaces for Router PE1:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.19.1.2/30;
      }
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.19.2.1/30;
      }
      family mpls;
    }
  }
}
```



```
}
```

## Configuring Routing Options on Router PE1

At the **[edit routing-options]** hierarchy level, you need to configure the **router-id** and **autonomous-system** statements. The **router-id** statement identifies Router PE1.

Configure the routing options for Router PE1:

```
[edit]
routing-options {
  router-id 10.255.14.216;
  autonomous-system 69;
}
```

## Configuring Protocols on Router PE1

On Router PE1, you need to configure MPLS, BGP, OSPF, and LDP at the **[edit protocols]** hierarchy level:

```
[edit]
protocols {
  mpls {
    interface so-0/0/1.0;
  }
  bgp {
    group San-Francisco-Chicago {
      type internal;
      preference 10;
      local-address 10.255.14.216;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.224;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/0/1.0;
    }
  }
  ldp {
    interface so-0/0/1.0;
  }
}
```

## Configuring Policy Options on Router PE1

On Router PE1, you need to configure policies at the **[edit policy-options]** hierarchy level. These policies ensure that the CE routers in the Layer 3 VPN exchange routing information. In this example, Router CE1 in San Francisco exchanges routing information with Router CE2 in Chicago.

Configure the policy options on the PE1 router:

```

[edit]
policy-options {
  policy-statement vpn-import-VPN-A {
    term term1 {
      from {
        protocol bgp;
        community import-target-VPN-A;
      }
      then accept;
    }
    term term2 {
      then reject;
    }
  }
  policy-statement vpn-export-VPN-A {
    term term1 {
      from protocol ospf;
      then {
        community add export-target-VPN-A;
        accept;
      }
    }
    term term2 {
      then reject;
    }
  }
  community export-target-VPN-A members [target:10.255.14.216:11
  domain-id:192.0.2.1:0];
  community import-target-VPN-A members target:10.255.14.224:31;
}

```

## Configuring the Routing Instance on Router PE1

You need to configure a Layer 3 VPN routing instance on Router PE1. To indicate that the routing instance is for a Layer 3 VPN, add the **instance-type vrf** statement at the **[edit routing-instance *routing-instance-name*]** hierarchy level.

The **domain-id** statement is configured at the **[edit routing-instances routing-options protocols ospf]** hierarchy level. As shown in [Figure 16 on page 180](#), the routing instance on Router PE2 must share the same domain ID as the corresponding routing instance on Router PE1 so that routes from Router CE1 to Router CE2 and vice versa are distributed as Type 3 LSAs. If you configure different OSPF domain IDs in the routing instances for Router PE1 and Router PE2, the routes from each CE router will be distributed as Type 5 LSAs.

Configure the routing instance on Router PE1:

```

[edit]
routing-instances {
  VPN-A-San-Francisco-Chicago {
    instance-type vrf;
    interface so-0/0/0.0;
    route-distinguisher 10.255.14.216:11;
    vrf-import vpn-import-VPN-A;
    vrf-export vpn-export-VPN-A;
  }
}

```

```

routing-options {
  router-id 10.255.14.216;
  autonomous-system 69;
}
protocols {
  ospf {
    domain-id 192.0.2.1;
    export vpn-import-VPN-A;
    area 0.0.0.0 {
      interface so-0/0/0.0;
    }
  }
}
}

```

### Configuration Summary for Router PE1

<b>Configure Interfaces</b>	<pre> interfaces {   so-0/0/0 {     unit 0 {       family inet {         address 10.19.1.2/30;       }     }   }   so-0/0/1 {     unit 0 {       family inet {         address 10.19.2.1/30;       }       family mpls;     }   } } </pre>
<b>Configure Routing Options</b>	<pre> routing-options {   router-id 10.255.14.216;   autonomous-system 69; } </pre>
<b>Configure Protocols</b>	<pre> protocols {   mpls {     interface so-0/0/0.0;   }   bgp {     group San-Francisco-Chicago {       type internal;       preference 10;       local-address 10.255.14.216;       family inet-vpn {         unicast;       }     }     neighbor 10.255.14.224;   } } </pre>

```

    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/0/1.0;
    }
  }
  ldp {
    interface so-0/0/1.0;
  }
}

```

**Configure VPN Policy**

```

policy-options {
  policy-statement vpn-import-VPN-A {
    term term1 {
      from {
        protocol bgp;
        community import-target-VPN-A;
      }
      then accept;
    }
    term term2 {
      then reject;
    }
  }
  policy-statement vpn-export-VPN-A {
    term term1 {
      from protocol ospf;
      then {
        community add export-target-VPN-A;
        accept;
      }
    }
    term term2 {
      then reject;
    }
  }
  community export-target-VPN-B members [ target:10.255.14.216:11domain-id:192.0.2.1:0
];
  community import-target-VPN-B members target:10.255.14.224:31;
}

```

**Routing Instance for  
Layer 3 VPN**

```

routing-instances {
  VPN-A-San-Francisco-Chicago {
    instance-type vrf;
    interface so-0/0/0.0;
    route-distinguisher 10.255.14.216:11;
    vrf-import vpn-import-VPN-A;
    vrf-export vpn-export-VPN-A;
    routing-options {
      router-id 10.255.14.216;
      autonomous-system 69;
    }
  }
}

```

```

protocols {
  ospf {
    domain-id 192.0.2.1;
    export vpn-import-VPN-A;
    area 0.0.0.0 {
      interface so-0/0/0.0;
    }
  }
}

```

### Example: Giving VPN Backbones Priority with OSPFv2 Sham Links for Layer 3 VPNs

- [OSPFv2 Sham Links Overview on page 185](#)
- [Example: Configuring OSPFv2 Sham Links on page 186](#)

#### OSPFv2 Sham Links Overview


You can create an intra-area link or sham link between two provider edge (PE) routing devices so that the VPN backbone is preferred over the back-door link. A back-door link is a backup link that connects customer edge (CE) devices in case the VPN backbone is unavailable. When such a backup link is available and the CE devices are in the same OSPF area, the default behavior is to prefer this backup link over the VPN backbone. This is because the backup link is considered an intra-area link, while the VPN backbone is always considered an interarea link. Intra-area links are always preferred over interarea links.

The sham link is an unnumbered point-to-point intra-area link between PE devices. When the VPN backbone has a sham intra-area link, this sham link can be preferred over the backup link if the sham link has a lower OSPF metric than the backup link.

The sham link is advertised using Type 1 link-state advertisements (LSAs). Sham links are valid only for routing instances and OSPFv2.

Each sham link is identified by the combination of a local endpoint address and a remote endpoint address. [Figure 17 on page 185](#) shows an OSPFv2 sham link. Router CE1 and Router CE2 are located in the same OSPFv2 area. These customer edge (CE) routing devices are linked together by a Layer 3 VPN over Router PE1 and Router PE2. In addition, Router CE1 and Router CE2 are connected by an intra-area link used as a backup.

*Figure 17: OSPFv2 Sham Link*



OSPFv2 treats the link through the Layer 3 VPN as an interarea link. By default, OSPFv2 prefers intra-area links to interarea links, so OSPFv2 selects the backup intra-area link as the active path. This is not acceptable in a configuration where the intra-area link is not the expected primary path for traffic between the CE routing devices. You can configure the metric for the sham link to ensure that the path over the Layer 3 VPN is preferred to a backup path over an intra-area link connecting the CE routing devices.

For the remote endpoint, you can configure the OSPFv2 interface as a demand circuit, configure IPsec authentication (you configure the actual IPsec authentication separately), and define the metric value.

You should configure an OSPFv2 sham link under the following circumstances:

- Two CE routing devices are linked together by a Layer 3 VPN.
- These CE routing devices are in the same OSPFv2 area.
- An intra-area link is configured between the two CE routing devices.

If there is no intra-area link between the CE routing devices, you do not need to configure an OSPFv2 sham link.



**NOTE:** In Junos OS Release 9.6 and later, an OSPFv2 sham link is installed in the routing table as a hidden route. Additionally, a BGP route is not exported to OSPFv2 if a corresponding OSPF sham link is available.

---



**NOTE:** In Junos OS Release 16.1 and later, OSPF sham-links are supported on default instances. The cost of the sham-link is dynamically set to the aigp-metric of the BGP route if no metric is configured on the sham-link by the user. If the aigp-metric is not present in the BGP route then the sham-link cost defaults to 1.

---

- See Also**
- [Example: Configuring OSPFv2 Sham Links on page 186](#)
  - For more information about Layer 3 VPNs, see the *Junos OS VPNs Library for Routing Devices*

## Example: Configuring OSPFv2 Sham Links

This example shows how to enable OSPFv2 sham links on a PE routing device.

- [Requirements on page 187](#)
- [Overview on page 187](#)
- [Configuration on page 188](#)
- [Verification on page 193](#)

## Requirements

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

## Overview

The sham link is an unnumbered point-to-point intra-area link and is advertised by means of a type 1 link-state advertisement (LSA). Sham links are valid only for routing instances and OSPFv2.


Each sham link is identified by a combination of the local endpoint address and a remote endpoint address and the OSPFv2 area to which it belongs. You manually configure the sham link between two PE devices, both of which are within the same VPN routing and forwarding (VRF) routing instance, and you specify the address for the local end point of the sham link. This address is used as the source for the sham link packets and is also used by the remote PE routing device as the sham link remote end point. You can also include the optional **metric** option to set a metric value for the remote end point. The metric value specifies the cost of using the link. Routes with lower total path metrics are preferred over those with higher path metrics.

To enable OSPFv2 sham links on a PE routing device:

- Configure an extra loopback interface on the PE routing device.
- Configure the VRF routing instance that supports Layer 3 VPNs on the PE routing device, and associate the sham link with an existing OSPF area. The OSPFv2 sham link configuration is also included in the routing instance. You configure the sham link's local endpoint address, which is the loopback address of the local VPN, and the remote endpoint address, which is the loopback address of the remote VPN. In this example, the VRF routing instance is named **red**.

Figure 18 on page 187 shows an OSPFv2 sham link.

Figure 18: OSPFv2 Sham Link Example



The devices in the figure represent the following functions:

- CE1 and CE2 are the customer edge devices.
- PE1 and PE2 are the provider edge devices.
- P is the provider device.

“CLI Quick Configuration” on page 188 shows the configuration for all of the devices in Figure 18 on page 187. The section “Step-by-Step Procedure” on page 190 describes the steps on Device PE1.

### 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.

```
CE1  set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.1/30
      set interfaces fe-1/2/0 unit 0 family mpls
      set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
      set interfaces lo0 unit 0 family inet address 192.0.2.1/24
      set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
      set protocols ospf area 0.0.0.0 interface lo0.0 passive
      set protocols ospf area 0.0.0.0 interface fe-1/2/1.0 metric 100
      set policy-options policy-statement send-direct from protocol direct
      set policy-options policy-statement send-direct then accept
      set routing-options router-id 192.0.2.1
      set routing-options autonomous-system 1

PE1  set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.2/30
      set interfaces fe-1/2/0 unit 0 family mpls
      set interfaces fe-1/2/1 unit 0 family inet address 10.1.1.5/30
      set interfaces fe-1/2/1 unit 0 family mpls
      set interfaces lo0 unit 0 family inet address 192.0.2.2/24
      set interfaces lo0 unit 1 family inet address 198.51.100.2/24
      set protocols mpls interface fe-1/2/1.0
      set protocols bgp group toR4 type internal
      set protocols bgp group toR4 local-address 192.0.2.2
      set protocols bgp group toR4 family inet-vpn unicast
      set protocols bgp group toR4 neighbor 192.0.2.4
      set protocols ospf area 0.0.0.0 interface fe-1/2/1.0
      set protocols ospf area 0.0.0.0 interface lo0.0 passive
      set protocols ldp interface fe-1/2/1.0
      set protocols ldp interface lo0.0
      set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
      set policy-options policy-statement bgp-to-ospf term 1 then accept
      set policy-options policy-statement bgp-to-ospf term 2 then reject
      set routing-instances red instance-type vrf
      set routing-instances red interface fe-1/2/0.0
      set routing-instances red interface lo0.1
      set routing-instances red route-distinguisher 2:1
      set routing-instances red vrf-target target:2:1
      set routing-instances red protocols ospf export bgp-to-ospf
```



```

set routing-instances red protocols ospf sham-link local 198.51.100.2
set routing-instances red protocols ospf area 0.0.0.0 sham-link-remote 198.51.100.4
  metric 10
set routing-instances red protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set routing-instances red protocols ospf area 0.0.0.0 interface lo0.1
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 2

P    set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.6/30
    set interfaces fe-1/2/0 unit 0 family mpls
    set interfaces fe-1/2/1 unit 0 family inet address 10.1.1.9/30
    set interfaces fe-1/2/1 unit 0 family mpls
    set interfaces lo0 unit 3 family inet address 192.0.2.3/24
    set protocols mpls interface all
    set protocols ospf area 0.0.0.0 interface lo0.3 passive
    set protocols ospf area 0.0.0.0 interface all
    set protocols ldp interface all
    set routing-options router-id 192.0.2.3

PE2  set interfaces fe-1/2/0 unit 0 family inet address 10.1.1.10/30
    set interfaces fe-1/2/0 unit 0 family mpls
    set interfaces fe-1/2/1 unit 0 family inet address 10.1.1.13/30
    set interfaces fe-1/2/1 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 192.0.2.4/32
    set interfaces lo0 unit 1 family inet address 198.51.100.4/32
    set protocols mpls interface fe-1/2/0.0
    set protocols bgp group toR2 type internal
    set protocols bgp group toR2 local-address 192.0.2.4
    set protocols bgp group toR2 family inet-vpn unicast
    set protocols bgp group toR2 neighbor 192.0.2.2
    set protocols ospf area 0.0.0.0 interface lo0.0 passive
    set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
    set protocols ldp interface fe-1/2/0.0
    set protocols ldp interface lo0.0
    set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
    set policy-options policy-statement bgp-to-ospf term 1 then accept
    set policy-options policy-statement bgp-to-ospf term 2 then reject
    set routing-instances red instance-type vrf
    set routing-instances red interface fe-1/2/1.0
    set routing-instances red interface lo0.1
    set routing-instances red route-distinguisher 2:1
    set routing-instances red vrf-target target:2:1
    set routing-instances red protocols ospf export bgp-to-ospf
    set routing-instances red protocols ospf sham-link local 198.51.100.4
    set routing-instances red protocols ospf area 0.0.0.0 sham-link-remote 198.51.100.2
      metric 10
    set routing-instances red protocols ospf area 0.0.0.0 interface fe-1/2/1.0
    set routing-instances red protocols ospf area 0.0.0.0 interface lo0.1
    set routing-options router-id 192.0.2.4
    set routing-options autonomous-system 2

CE2  set interfaces fe-1/2/0 unit 14 family inet address 10.1.1.14/30
    set interfaces fe-1/2/0 unit 14 family mpls
    set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30

```

```

set interfaces lo0 unit 5 family inet address 192.0.2.5/24
set protocols ospf area 0.0.0.0 interface fe-1/2/0.14
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.18
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 3

```

### Step-by-Step Procedure

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

To configure OSPFv2 sham links on each PE device:

1. Configure the interfaces, including two loopback interfaces.

```

[edit interfaces]
user@PE1# set fe-1/2/0 unit 0 family inet address 10.1.1.2/30
user@PE1# set fe-1/2/0 unit 0 family mpls
user@PE1# set fe-1/2/1 unit 0 family inet address 10.1.1.5/30
user@PE1# set fe-1/2/1 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 192.0.2.2/24
user@PE1# set lo0 unit 1 family inet address 198.51.100.2/24

```

2. Configure MPLS on the core-facing interface.

```

[edit protocols mpls]
user@PE1# set interface fe-1/2/1.0

```

3. Configure internal BGP (IBGP).

```

[edit ]
user@PE1# set protocols bgp group toR4 type internal
user@PE1# set protocols bgp group toR4 local-address 192.0.2.2
user@PE1# set protocols bgp group toR4 family inet-vpn unicast
user@PE1# set protocols bgp group toR4 neighbor 192.0.2.4

```

4. Configure OSPF on the core-facing interface and on the loopback interface that is being used in the main instance.

```

[edit protocols ospf area 0.0.0.0]
user@PE1# set interface fe-1/2/1.0
user@PE1# set interface lo0.0 passive

```

5. Configure LDP or RSVP on the core-facing interface and on the loopback interface that is being used in the main instance.

```

[edit protocols ldp]
user@PE1# set interface fe-1/2/1.0
user@PE1# set interface lo0.0

```

6. Configure a routing policy for use in the routing instance.

```
[edit policy-options policy-statement bgp-to-ospf]
user@PE1# set term 1 from protocol bgp
user@PE1# set term 1 then accept
user@PE1# set term 2 then reject
```

7. Configure the routing instance.

```
[edit routing-instances red]
user@PE1# set instance-type vrf
user@PE1# set interface fe-1/2/0.0
user@PE1# set route-distinguisher 2:1
user@PE1# set vrf-target target:2:1
user@PE1# set protocols ospf export bgp-to-ospf
user@PE1# set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
```

8. Configure the OSPFv2 sham link.

Include the extra loopback interface in the routing instance and also in the OSPF configuration.

Notice that the metric on the sham-link interface is set to 10. On Device CE1's backup OSPF link, the metric is set to 100. This causes the sham link to be the preferred link.

```
[edit routing-instances red]
user@PE1# set interface lo0.1
user@PE1# set protocols ospf sham-link local 198.51.100.2
user@PE1# set protocols ospf area 0.0.0.0 sham-link-remote 198.51.100.4 metric
10
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.1
```

9. Configure the autonomous system (AS) number and the router ID.

```
[edit routing-options]
user@PE1# set router-id 192.0.2.2
user@PE1# set autonomous-system 2
```

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

```
[edit]
user@R1# commit
```

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

Output for PE1:

```
user@PE1# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.1.1.2/30;
```

```
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.1.1.5/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
  unit 1 {
    family inet {
      address 198.51.100.2/24;
    }
  }
}

user@PE1# show protocols
mpls {
  interface fe-1/2/1.0;
}
bgp {
  group toR4 {
    type internal;
    local-address 192.0.2.2;
    family inet-vpn {
      unicast;
    }
    neighbor 192.0.2.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface fe-1/2/1.0;
  interface lo0.0;
}

user@PE1# show policy-options
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
}
```

```

    }
    term 2 {
        then reject;
    }
}

user@PE1# show routing-instances
red {
    instance-type vrf;
    interface fe-1/2/0.0;
    interface lo0.1;
    route-distinguisher 2:1;
    vrf-target target:2:1;
    protocols {
        ospf {
            export bgp-to-ospf;
            sham-link local 198.51.100.2;
            area 0.0.0.0 {
                sham-link-remote 198.51.100.4 metric 10;
                interface fe-1/2/0.0;
                interface lo0.1;
            }
        }
    }
}

user@PE1# show routing-options
router-id 192.0.2.2;
autonomous-system 2;

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Sham Link Interfaces on page 193](#)
- [Verifying the Local and Remote End Points of the Sham Link on page 194](#)
- [Verifying the Sham Link Adjacencies on page 194](#)
- [Verifying the Link-State Advertisement on page 194](#)
- [Verifying the Path Selection on page 195](#)

### Verifying the Sham Link Interfaces

**Purpose** Verify the sham link interface. The sham link is treated as an interface in OSPFv2, with the named displayed as **shamlink.<unique identifier>**, where the unique identifier is a number. For example, **shamlink.0**. The sham link appears as a point-to-point interface.

**Action** From operational mode, enter the **show ospf interface instance *instance-name*** command.

```

user@PE1> show ospf interface instance red

```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.1	DR	0.0.0.0	198.51.100.2	0.0.0.0	
0					
fe-1/2/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
shamlink.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

**Verifying the Local and Remote End Points of the Sham Link**

**Purpose** Verify the local and remote end points of the sham link. The MTU for the sham link interface is always zero.

**Action** From operational mode, enter the **show ospf interface instance *instance-name* detail** command.

```
user@PE1> show ospf interface shamlink.0 instance red
Interface      State   Area      DR ID      BDR ID      Nbrs
shamlink.0     PtToPt  0.0.0.0    0.0.0.0    0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 0, Cost: 10
Local: 198.51.100.2, Remote: 198.51.100.4
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Protection type: None, No eligible backup
Topology default (ID 0) -> Cost: 10
```

**Verifying the Sham Link Adjacencies**

**Purpose** Verify the adjacencies between the configured sham links.

**Action** From operational mode, enter the **show ospf neighbor instance *instance-name*** command.

```
user@PE1> show ospf neighbor instance red
Address        Interface      State   ID          Pri  Dead
10.1.1.1       fe-1/2/0.0    Full   192.0.2.1   128   35
198.51.100.4   shamlink.0    Full   198.51.100.4 0
31
```

**Verifying the Link-State Advertisement**

**Purpose** Verify that the router LSA originated by the instance carries the sham link adjacency as an unnumbered point-to-point link. The link data for sham links is a number ranging from 0x80010000 through 0x8001ffff.

**Action** From operational mode, enter the **show ospf database instance *instance-name*** command.

```
user@PE1> show ospf database instance red

OSPF database, Area 0.0.0.0
Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router    192.0.2.1    192.0.2.1    0x80000009 1803 0x22 0x6ec7 72
Router    192.0.2.5    192.0.2.5    0x80000007 70   0x22 0x2746 72
Router    *198.51.100.2 198.51.100.2 0x80000006 55   0x22 0xda6b 60
Router    198.51.100.4 198.51.100.4 0x80000005 63   0x22 0xb19 60
```

```

Network 10.0.0.18      192.0.2.5      0x80000002    70 0x22 0x9a71 32
  OSPF AS SCOPE link state database
  Type      ID      Adv Rtr      Seq      Age Opt Cksum Len
Extern 198.51.100.2      198.51.100.4      0x80000002    72 0xa2 0x343
36
Extern *198.51.100.4      198.51.100.2      0x80000002    71 0xa2 0xe263
36

```

### Verifying the Path Selection

**Purpose** Verify that the Layer 3 VPN path is used instead of the backup path.

**Action** From operational mode, enter the **traceroute** command from Device CE1 to Device CE2.

```

user@CE1> traceroute 192.0.2.5

traceroute to 192.0.2.5 (192.0.2.5), 30 hops max, 40 byte packets
 1 10.1.1.2 (10.1.1.2) 1.930 ms 1.664 ms 1.643 ms
 2 * * *
 3 10.1.1.10 (10.1.1.10) 2.485 ms 1.435 ms 1.422 ms
   MPLS Label=299808 CoS=0 TTL=1 S=1
 4 192.0.2.5 (192.0.2.5) 1.347 ms 1.362 ms 1.329 ms

```

**Meaning** The traceroute operation shows that the Layer 3 VPN is the preferred path. If you were to remove the sham link or if you were to modify the OSPF metric to prefer that backup path, the traceroute would show that the backup path is preferred.

**See Also** • [OSPFv2 Sham Links Overview on page 185](#)

**Related Documentation** • [Understanding OSPF Configurations](#)  
• [Junos OS VPNs Library for Routing Devices](#)

## Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs

You can configure an EBGp or IBGP multihop session between the PE and CE routers of a Layer 3 VPN. This allows you to have one or more routers between the PE and CE routers. Using IBGP between PE and CE routers does not require the configuration of any additional statements. However, using EBGp between the PE and CE routers requires the configuration of the **multihop** statement.

To configure an external BGP multihop session for the connection between the PE and CE routers, include the **multihop** statement on the PE router. To help prevent routing loops, you have to configure a time-to-live (TTL) value for the multihop session:

```
multihop ttl-value;
```


For the list of hierarchy levels at which you can configure this statement, see the summary section for this statement.

## Configuring an LDP-over-RSVP VPN Topology

This example shows how to set up a VPN topology in which LDP packets are tunneled over an RSVP LSP. This configuration consists of the following components (see [Figure 19 on page 196](#)):

- One VPN (VPN-A)
- Two PE routers
- LDP as the signaling protocol between the PE routers and their adjacent P routers
- An RSVP LSP between two of the P routers over which LDP is tunneled

*Figure 19: Example of an LDP-over-RSVP VPN Topology*



The following steps describe how this topology is established and how packets are sent from CE Router CE2 to CE Router CE1:

1. The P routers P1 and P3 establish RSVP LSPs between each other and install their loopback addresses in their inet.3 routing tables.
2. PE Router PE1 establishes an LDP session with Router P1 over interface **so-1/0/0.0**.
3. Router P1 establishes an LDP session with Router P3's loopback address, which is reachable using the RSVP LSP.
4. Router P1 sends its label bindings, which include a label to reach Router PE1, to Router P3. These label bindings allow Router P3 to direct LDP packets to Router PE1.
5. Router P3 establishes an LDP session with Router PE2 over interface **so0-0/0/0.0** and establishes an LDP session with Router P1's loopback address.
6. Router P3 sends its label bindings, which include a label to reach Router PE2, to Router P1. These label bindings allow Router P1 to direct LDP packets to Router PE2's loopback address.



7. Routers PE1 and PE2 establish IBGP sessions with each other.
8. When Router PE1 announces to Router PE2 routes that it learned from Router CE1, it includes its VPN label. (The PE router creates the VPN label and binds it to the interface between the PE and CE routers.) Similarly, when Router PE2 announces routes that it learned from Router CE2, it sends its VPN label to Router PE1.

When Router PE2 wants to forward a packet to Router CE1, it pushes two labels onto the packet's label stack: first the VPN label that is bound to the interface between Router PE1 and Router CE1, then the LDP label used to reach Router PE1. Then it forwards the packets to Router P3 over interface **so-0/0/1.0**.

1. When Router P3 receives the packets from Router PE2, it swaps the LDP label that is on top of the stack (according to its LDP database) and also pushes an RSVP label onto the top of the stack so that the packet can now be switched by the RSVP LSP. At this point, there are three labels on the stack: the inner (bottom) label is the VPN label, the middle is the LDP label, and the outer (top) is the RSVP label.
2. Router P2 receives the packet and switches it to Router P1 by swapping the RSVP label. In this topology, because Router P2 is the penultimate-hop router in the LSP, it pops the RSVP label and forwards the packet over interface **so-1/1/0.0** to Router P1. At this point, there are two labels on the stack: The inner label is the VPN label, and the outer one is the LDP label.
3. When Router P1 receives the packet, it pops the outer label (the LDP label) and forwards the packet to Router PE1 using interface **so-1/0/0.0**. In this topology, Router PE1 is the egress LDP router, so Router P1 pops the LDP label instead of swapping it with another label. At this point, there is only one label on the stack, the VPN label.
4. When Router PE1 receives the packet, it pops the VPN label and forwards the packet as an IPv4 packet to Router CE1 over interface **ge-1/1/0.0**.

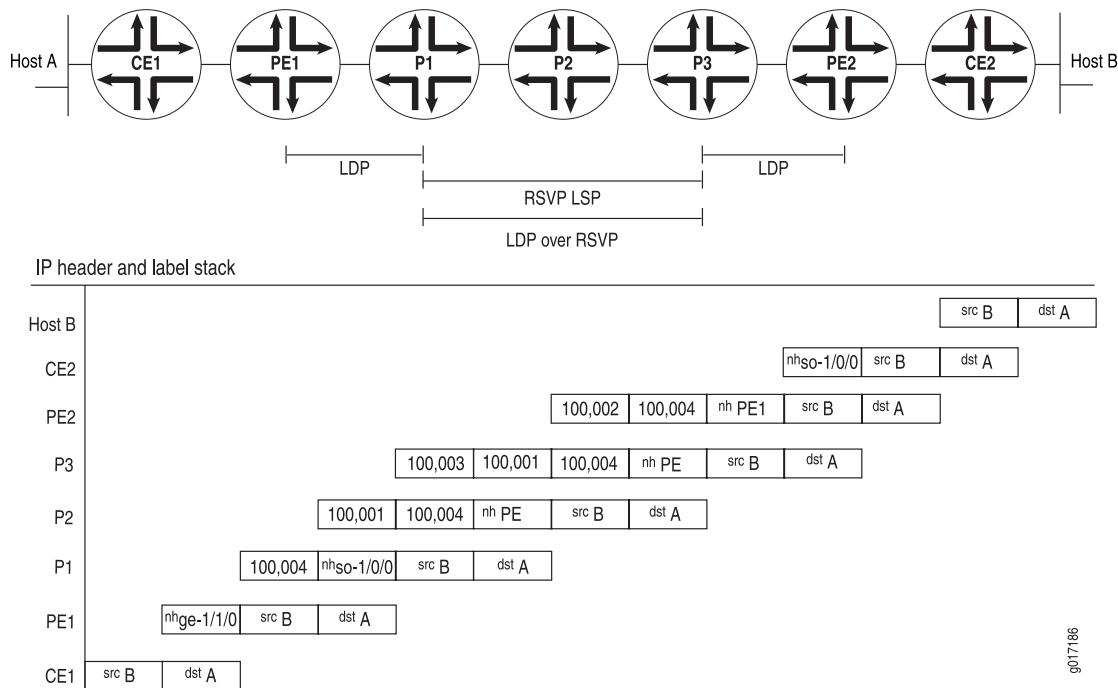
A similar set of operations occurs for packets sent from Router CE1 that are destined for Router CE2.

The following list explains how, for packets being sent from Router CE2 to Router CE1, the LDP, RSVP, and VPN labels are announced by the various routers. These steps include examples of label values (illustrated in [Figure 20 on page 198](#)).

- LDP labels
  - Router PE1 announces LDP label 3 for itself to Router P1.
  - Router P1 announces LDP label 100,001 for Router PE1 to Router P3.
  - Router P3 announces LDP label 100,002 for Router PE1 to Router PE2.
- RSVP labels
  - Router P1 announces RSVP label 3 to Router P2.
  - Router P2 announces RSVP label 100,003 to Router P3.

- VPN label
- Router PE1 announces VPN label 100,004 to Router PE2 for the route from Router CE1 to Router CE2.

Figure 20: Label Pushing and Popping



For a packet sent from Host B in [Figure 20 on page 198](#) to Host A, the packet headers and labels change as the packet travels to its destination:

1. The packet that originates from Host B has a source address of B and a destination address of A in its header.
2. Router CE2 adds to the packet a next hop of interface **so-1/0/0**.
3. Router PE2 swaps out the next hop of interface **so-1/0/0** and replaces it with a next hop of PE1. It also adds two labels for reaching Router PE1, first the VPN label (100,004), then the LDP label (100,002). The VPN label is thus the inner (bottom) label on the stack, and the LDP label is the outer label.
4. Router P3 swaps out the LDP label added by Router PE2 (100,002) and replaces it with its LDP label for reaching Router PE1 (100,001). It also adds the RSVP label for reaching Router P2 (100,003).
5. Router P2 removes the RSVP label (100,003) because it is the penultimate hop in the MPLS LSP.
6. Router P1 removes the LDP label (100,001) because it is the penultimate LDP router. It also swaps out the next hop of PE1 and replaces it with the next-hop interface, **so-1/0/0**.

7. Router PE1 removes the VPN label (100,004). It also swaps out the next-hop interface of **so-1/0/0** and replaces it with its next-hop interface, **ge-1/1/0**.
8. Router CE1 removes the next-hop interface of **ge-1/1/0**, and the packet header now contains just a source address of B and a destination address of A.

The final section in this example consolidates the statements needed to configure VPN functionality on each of the service P routers shown in [Figure 19 on page 196](#).



**NOTE:** In this example, a private AS number is used for the route distinguisher and the route target. This number is used for illustration only. When you are configuring VPNs, you should use an assigned AS number.

The following sections explain how to configure the VPN functionality on the PE and P routers. The CE routers do not have any information about the VPN, so you configure them normally.

- [Enabling an IGP on the PE and P Routers on page 199](#)
- [Enabling LDP on the PE and P Routers on page 199](#)
- [Enabling RSVP and MPLS on the P Router on page 201](#)
- [Configuring the MPLS LSP Tunnel Between the P Routers on page 201](#)
- [Configuring IBGP on the PE Routers on page 202](#)
- [Configuring Routing Instances for VPNs on the PE Routers on page 203](#)
- [Configuring VPN Policy on the PE Routers on page 204](#)
- [LDP-over-RSVP VPN Configuration Summarized by Router on page 206](#)

## Enabling an IGP on the PE and P Routers

To allow the PE and P routers to exchange routing information among themselves, you must configure an IGP on all these routers or you must configure static routes. You configure the IGP on the master instance of the routing protocol process (rpd) (that is, at the **[edit protocols]** hierarchy level), not within the VPN routing instance (that is, not at the **[edit routing-instances]** hierarchy level).

You configure the IGP in the standard way. This configuration example does not include this portion of the configuration.

## Enabling LDP on the PE and P Routers

In this configuration example, the LDP is the signaling protocol between the PE routers. For the VPN to function, you must configure LDP on the two PE routers and on the P routers that are connected to the PE routers. You need to configure LDP only on the interfaces in the core of the service provider's network; that is, between the PE and P routers and between the P routers. You do not need to configure LDP on the interface between the PE and CE routers.

In this configuration example, you configure LDP on the P routers' loopback interfaces because these are the interfaces on which the MPLS LSP is configured.

On the PE routers, you must also configure **family inet** when you configure the logical interface.

On Router PE1, configure LDP:

```
[edit protocols]
ldp {
  interface so-1/0/0.0;
}
[edit interfaces]
so-1/0/0 {
  unit 0 {
    family mpls;
  }
}
```

On Router PE2, configure LDP:

```
[edit protocols]
ldp {
  interface so-0/0/0.0;
}
[edit interfaces]
so-0/0/1 {
  unit 0 {
    family mpls;
  }
}
```

On Router P1, configure LDP:

```
[edit protocols]
ldp {
  interface so-1/0/0.0;
  interface lo0;
}
```

On Router P3, configure LDP:

```
[edit protocols]
ldp {
  interface lo0;
  interface so-0/0/0.0;
}
```

On Router P2, although you do not need to configure LDP, you can optionally configure it to provide a fallback LDP path in case the RSVP LSP becomes nonoperational:

```
[edit protocols]
ldp {
  interface so-1/1/0.0;
  interface at-2/0/0.0;
}
```

## Enabling RSVP and MPLS on the P Router

On the P Router P2 you must configure RSVP and MPLS because this router exists on the MPLS LSP path between the P Routers P1 and P3:

```
[edit]
protocols {
  rsvp {
    interface so-1/1/0.0;
    interface at-2/0/0.0;
  }
  mpls {
    interface so-1/1/0.0;
    interface at-2/0/0.0;
  }
}
```

## Configuring the MPLS LSP Tunnel Between the P Routers

In this configuration example, LDP is tunneled over an RSVP LSP. Therefore, in addition to configuring RSVP, you must enable traffic engineering support in an IGP, and you must create an MPLS LSP to tunnel the LDP traffic.

On Router P1, enable RSVP and configure one end of the MPLS LSP tunnel. In this example, traffic engineering support is enabled for OSPF, and you configure MPLS on the interfaces to the LSP and to Router PE1. In the **to** statement, you specify the loopback address of Router P3.

```
[edit]
protocols {
  rsvp {
    interface so-1/0/1.0;
  }
  mpls {
    label-switched-path P1-to-P3 {
      to 10.255.100.1;
      ldp-tunneling;
    }
    interface so-1/0/0.0;
    interface so-1/0/1.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-1/0/0.0;
      interface so-1/0/1.0;
    }
  }
}
```

On Router P3, enable RSVP and configure the other end of the MPLS LSP tunnel. Again, traffic engineering support is enabled for OSPF, and you configure MPLS on the interfaces to the LSP and to Router PE2. In the **to** statement, you specify the loopback address of Router P1.

```
[edit]
protocols {
  rsvp {
    interface at-2/0/1.0;
  }
  mpls {
    label-switched-path P3-to-P1 {
      to 10.255.2.2;
      ldp-tunneling;
    }
    interface at-2/0/1.0;
    interface so-0/0/0.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface at-2/0/1.0;
      interface so-0/0/0.0;
    }
  }
}
```

## Configuring IBGP on the PE Routers

On the PE routers, configure an IBGP session with the following properties:

- VPN family—To indicate that the IBGP session is for the VPN, include the **family inet-vpn** statement.
- Loopback address—Include the **local-address** statement, specifying the local PE router's loopback address. The IBGP session for VPNs runs through the loopback address. You must also configure the **lo0** interface at the **[edit interfaces]** hierarchy level. The example does not include this part of the router's configuration.
- Neighbor address—Include the **neighbor** statement, specifying the IP address of the neighboring PE router, which is its loopback (**lo0**) address.

On Router PE1, configure IBGP:

```
[edit]
protocols {
  bgp {
    group PE1-to-PE2 {
      type internal;
      local-address 10.255.1.1;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.200.2;
    }
  }
}
```

On Router PE2, configure IBGP:

```
[edit]
```

```

protocols {
  bgp {
    group PE2-to-PE1 {
      type internal;
      local-address 10.255.200.2;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.1.1;
    }
  }
}

```

### Configuring Routing Instances for VPNs on the PE Routers

Both PE routers service VPN-A, so you must configure one routing instance on each router for the VPN in which you define the following:

- Route distinguisher, which must be unique for each routing instance on the PE router. It is used to distinguish the addresses in one VPN from those in another VPN.
- Instance type of **vrf**, which creates the VRF table on the PE router.
- Interfaces connected to the CE routers.
- VRF import and export policies, which must be the same on each PE router that services the same VPN. Unless the import policy contains only a **then reject** statement, it must include reference to a community. Otherwise, when you try to commit the configuration, the commit fails.



**NOTE:** In this example, a private AS number is used for the route distinguisher. This number is used for illustration only. When you are configuring VPNs, you should use an assigned AS number.

- Routing between the PE and CE routers, which is required for the PE router to distribute VPN-related routes to and from connected CE routers. You can configure a routing protocol—BGP, OSPF, or RIP—or you can configure static routing.

On Router PE1, configure the following routing instance for VPN-A. In this example, Router PE1 uses RIP to distribute routes to and from the CE router to which it is connected.

```

[edit]
routing-instance {
  VPN-A {
    instance-type vrf;
    interface ge-1/0/0.0;
    route-distinguisher 65535:0;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    protocols {
      rip {
        group PE1-to-CE1 {
          neighbor ge-1/0/0.0;
        }
      }
    }
  }
}

```

```

    }
  }
}

```

On Router PE2, configure the following routing instance for VPN-A. In this example, Router PE2 uses OSPF to distribute routes to and from the CE router to which it is connected.

```

[edit]
routing-instance {
  VPN-A {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 65535:1;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    protocols {
      ospf {
        area 0.0.0.0 {
          interface so-1/2/0.0;
        }
      }
    }
  }
}

```

## Configuring VPN Policy on the PE Routers

You must configure VPN import and export policies on each of the PE routers so that they install the appropriate routes in their VRF tables, which they use to forward packets within a VPN. For VPN-A, the VRF table is VPN-A.inet.0.

In the VPN policy, you also configure VPN target communities.



**NOTE:** In this example, a private AS number is used for the route target. This number is used for illustration only. When you are configuring VPNs, you should use an assigned AS number.

On Router PE1, configure the following VPN import and export policies:



**NOTE:** The policy qualifiers shown in this example are only those needed for the VPN to function. You can configure additional qualifiers, as needed, to any policies that you configure.

```

[edit]
policy-options {
  policy-statement VPN-A-import {
    term a {
      from {
        protocol bgp;

```



```

        community VPN-A;
    }
    then accept;
}
term b {
    then reject;
}
}
policy-statement VPN-A-export {
    term a {
        from protocol rip;
        then {
            community add VPN-A;
            accept;
        }
    }
    term b {
        then reject;
    }
}
community VPN-A members target:65535:00;
}

```

On Router PE2, configure the following VPN import and export policies:

```

[edit]
policy-options {
    policy-statement VPN-A-import {
        term a {
            from {
                protocol bgp;
                community VPN-A;
            }
            then accept;
        }
        term b {
            then reject;
        }
    }
    policy-statement VPN-A-export {
        term a {
            from protocol ospf;
            then {
                community add VPN-A;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
    community VPN-A members target:65535:00;
}

```

To apply the VPN policies on the routers, include the **vrf-export** and **vrf-import** statements when you configure the routing instance on the PE routers. The VRF import and export

policies handle the route distribution across the IBGP session running between the PE routers.

## LDP-over-RSVP VPN Configuration Summarized by Router

### Router PE1

<b>Routing Instance for VPN-A</b>	<pre> routing-instance {   VPN-A {     instance-type vrf;     interface ge-1/0/0.0;     route-distinguisher 65535:0;     vrf-import VPN-A-import;     vrf-export VPN-A-export;   } } </pre>
<b>Instance Routing Protocol</b>	<pre> protocols {   rip {     group PE1-to-CE1 {       neighbor ge-1/0/0.0;     }   } } </pre>
<b>Interfaces</b>	<pre> interfaces {   so-1/0/0 {     unit 0 {       family mpls;     }   }   ge-1/0/0 {     unit 0;   } } </pre>
<b>Master Protocol Instance</b>	<pre> protocols { } </pre>
<b>Enable LDP</b>	<pre> ldp {   interface so-1/0/0.0; } </pre>
<b>Enable MPLS</b>	<pre> mpls {   interface so-1/0/0.0;   interface ge-1/0/0.0; } </pre>
<b>Configure IBGP</b>	<pre> bgp {   group PE1-to-PE2 {     type internal;     local-address 10.255.1.1;   } } </pre>

```

    family inet-vpn {
        unicast;
    }
    neighbor 10.255.100.1;
}

```

**Configure VPN Policy**

```

policy-options {
    policy-statement VPN-A-import {
        term a {
            from {
                protocol bgp;
                community VPN-A;
            }
            then accept;
        }
        term b {
            then reject;
        }
    }
    policy-statement VPN-A-export {
        term a {
            from protocol rip;
            then {
                community add VPN-A;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
    community VPN-A members target:65535:00;
}

```

**Router P1**

<b>Master Protocol Instance</b>	<pre>protocols { }</pre>
<b>Enable RSVP</b>	<pre>rsvp {     interface so-1/0/1.0; }</pre>
<b>Enable LDP</b>	<pre>ldp {     interface so-1/0/0.0;     interface lo0.0; }</pre>
<b>Enable MPLS</b>	<pre>mpls {     label-switched-path P1-to-P3 {         to 10.255.100.1;     } }</pre>

	<pre>        ldp-tunneling;       }       interface so-1/0/0.0;       interface so-1/0/1.0;     }   }  Configure OSPF for Traffic Engineering Support    ospf {             traffic-engineering;             area 0.0.0.0 {               interface so-1/0/0.0;               interface so-1/0/1.0;             }           }         }</pre>
--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

---

### Router P2

Master Protocol Instance	<pre>protocols { }</pre>
Enable RSVP	<pre>rsvp {   interface so-1/1/0.0;   interface at-2/0/0.0; }</pre>
Enable MPLS	<pre>mpls {   interface so-1/1/0.0;   interface at-2/0/0.0; }</pre>

---

### Router P3

Master Protocol Instance	<pre>protocols { }</pre>
Enable RSVP	<pre>rsvp {   interface at-2/0/1.0; }</pre>
Enable LDP	<pre>ldp {   interface so-0/0/0.0;   interface lo0.0; }</pre>
Enable MPLS	<pre>mpls {   label-switched-path P3-to-P1 {     to 10.255.2.2;     ldp-tunneling;   }   interface at-2/0/1.0;   interface so-0/0/0.0; }</pre>

**Configure OSPF for  
Traffic Engineering  
Support**

```
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface at-2/0/1.0;
    interface at-2/0/1.0;
  }
}
```

**Router PE2****Routing Instance for  
VPN-A**

```
routing-instance {
  VPN-A {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 65535:1;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
  }
}
```

**Instance Routing  
Protocol**

```
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-1/2/0.0;
    }
  }
}
```

**Interfaces**

```
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
  so-1/2/0 {
    unit 0;
  }
}
```

**Master Protocol  
Instance**

```
protocols {
}
```

**Enable LDP**

```
ldp {
  interface so-0/0/0.0;
}
```

**Enable MPLS**

```
mpls {
  interface so-0/0/0.0;
  interface so-1/2/0.0;
}
```

**Configure IBGP**

```
bgp {  
  group PE2-to-PE1 {  
    type internal;  
    local-address 10.255.200.2;  
    family inet-vpn {  
      unicast;  
    }  
    neighbor 10.255.1.1;  
  }  
}
```

**Configure VPN Policy**

```
policy-options {  
  policy-statement VPN-A-import {  
    term a {  
      from {  
        protocol bgp;  
        community VPN-A;  
      }  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement VPN-A-export {  
    term a {  
      from protocol ospf;  
      then {  
        community add VPN-A;  
        accept;  
      }  
    }  
    term b {  
      then reject;  
    }  
  }  
  community VPN-A members target:65535:01;  
}
```

## Associating Interfaces with Layer 3 VPNs

- [Configuring an Application-Based Layer 3 VPN Topology on page 211](#)

### Configuring an Application-Based Layer 3 VPN Topology

---

This example illustrates an application-based mechanism for forwarding traffic into a Layer 3 VPN. Typically, one or more interfaces are associated with, or bound to, a VPN by including them in the configuration of the VPN routing instance. By binding the interface to the VPN, the VPN's VRF table is used to make forwarding decisions for any incoming traffic on that interface. Binding the interface also includes the interface local routes in the VRF table, which provides next-hop resolution for VRF routes.

In this example, a firewall filter is used to define which incoming traffic on an interface is forwarded by means of the standard routing table, `inet.0`, and which incoming traffic is forwarded by means of the VRF table. You can expand this example such that incoming traffic on an interface can be redirected to one or more VPNs. For example, you can define a configuration to support a VPN that forwards traffic based on source address, that forwards Hypertext Transfer Protocol (HTTP) traffic, or that forwards only streaming media.

For this configuration to work, the following conditions must be true:

- The interfaces that use filter-based forwarding must not be bound to the VPN.
- Static routing must be used as the means of routing.
- You must define an interface routing table group that is shared among `inet.0` and the VRF tables to provide local routes to the VRF table.

This example consists of two client hosts (Client D and Client E) that are in two different VPNs and that want to send traffic both within the VPN and to the Internet. The paths are defined as follows:

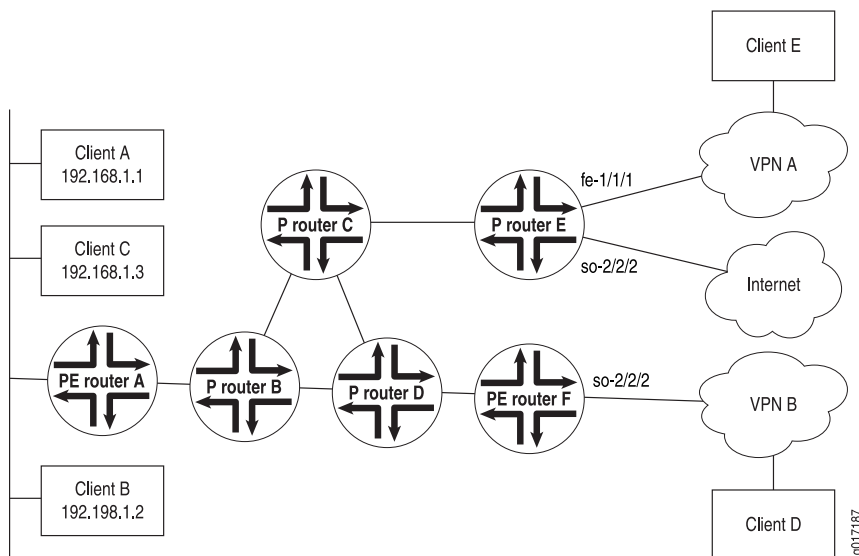
- Client A sends traffic to Client E over VPN A with a return path that also uses VPN A (using the VPN's VRF table).
- Client B sends traffic to Client D over VPN B with a return path that uses standard destination-based routing (using the `inet.0` routing table).
- Clients B and C send traffic to the Internet using standard routing (using the `inet.0` routing table), with a return path that also uses standard routing.

This example illustrates that there are a large variety of options in configuring an application-based Layer 3 VPN topology. This flexibility has application in many network implementations that require specific traffic to be forwarded in a constrained routing environment.

This configuration example shows only the portions of the configuration for the filter-based forwarding, routing instances, and policy. It does not illustrate how to configure a Layer 3 VPN.

[Figure 21 on page 212](#) illustrates the network topology used in this example.

**Figure 21: Application-Based Layer 3 VPN Example Configuration**



- [Configuration on Router A on page 212](#)
- [Configuration on Router E on page 214](#)
- [Configuration on Router F on page 214](#)

## Configuration on Router A

On Router A, you configure the interface to Clients A, B, and C. The configuration evaluates incoming traffic to determine whether it is to be forwarded by means of VPN or standard destination-based routing.

First, you apply an inbound filter and configure the interface:

```
[edit]
interfaces {
  fe-1/1/0 {
    unit 0 {
      family inet {
        filter {
          input fbf-vrf;
        }
        address 192.168.1.1/24;
      }
    }
  }
}
```



```

    }
  }
}

```

Because the interfaces that use filter-based forwarding must not be bound to a VPN, you must configure an alternate method to provide next-hop routes to the VRF table. You do this by defining an interface routing table group and sharing this group among all the routing tables:

```

[edit]
routing-options {
  interface-routes {
    rib-group inet if-rib;
  }
  rib-groups {
    if-rib {
      import-rib [ inet.0 vpn-A.inet.0 vpn-B.inet.0 ];
    }
  }
}

```

You apply the following filter to incoming traffic on interface **fe-1/1/0.0**. The first term matches traffic from Client A and forwards it to the routing instance for VPN A. The second term matches traffic from Client B that is destined for Client D and forwards it to the routing instance for VPN B. The third term matches all other traffic, which is forwarded normally by means of destination-based forwarding according to the routes in inet.0.

```

[edit firewall family family-name]
filter fbf-vrf {
  term vpnA {
    from {
      source-address {
        192.168.1.1/32;
      }
    }
    then {
      routing-instance vpn-A;
    }
  }
  term vpnB {
    from {
      source-address {
        192.168.1.2/32;
      }
      destination-address {
        192.168.3.0/24;
      }
    }
    then routing-instance vpn-B;
  }
}
term internet {
  then accept;
}

```

You then configure the routing instances for VPN A and VPN B. Notice that these statements include all the required statements to define a Layer 3 VPN except for the **interface** statement.

```
[edit]
routing-instances {
  vpn-A {
    instance-type vrf;
    route-distinguisher 172.21.10.63:100;
    vrf-import vpn-A-import;
    vrf-export vpn-A-export;
  }
  vpn-B {
    instance-type vrf;
    route-distinguisher 172.21.10.63:200;
    vrf-import vpn-B-import;
    vrf-export vpn-B-export;
  }
}
```

## Configuration on Router E

On Router E, configure a default route to reach the Internet. You should inject this route into the local IBGP mesh to provide an exit point from the network.

```
[edit]
routing-options {
  static {
    route 0.0.0.0/0 next-hop so-2/2/2.0 discard
  }
}
```

Configure the interface to Client E so that all incoming traffic on interface **fe-1/1/1.0** that matches the VPN policy is forwarded over VPN A:

```
[edit]
routing-instances {
  vpn-A {
    interface fe-1/1/1.0
    instance-type vrf;
    route-distinguisher 172.21.10.62:100;
    vrf-import vpn-A-import;
    vrf-export vpn-A-export;
    routing-options {
      static {
        route 192.168.2.0/24 next-hop fe-1/1/1.0;
      }
    }
  }
}
```

## Configuration on Router F

Again, because the interfaces that use filter-based forwarding must not be bound to a VPN, you configure an alternate method to provide next-hop routes to the VRF table by defining an interface routing table group and sharing this group among all the routing

tables. To provide a route back to the clients for normal inet.0 routing, you define a static route to include in inet.0 and redistribute the static route into BGP:

```
[edit]
routing-options {
  interface-routes {
    rib-group inet if-rib;
  }
  rib-groups {
    if-rib {
      import-rib [ inet.0 vpn-B.inet.0 ];
    }
  }
}
```

To direct traffic from VPN B to Client D, you configure the routing instance for VPN B on Router F. All incoming traffic from Client D on interface **so-3/3/3.0** is forwarded normally by means of the destination address based on the routes in inet.0.

```
[edit]
routing-instances {
  vpn-B {
    instance-type vrf;
    route-distinguisher 172.21.10.64:200;
    vrf-import vpn-B-import;
    vrf-export vpn-B-export;
    routing-options {
      static {
        route 192.168.3.0/24 next-hop so-3/3/3.0;
      }
    }
  }
}
```



# Carrying IPv4 Traffic Over Layer 3 VPNs


- [Understanding VPN-IPv4 Addresses and Route Distinguishers on page 217](#)
- [Understanding IPv4 Route Distribution in a Layer 3 VPN on page 221](#)
- [Configuring IPv4 Packet Forwarding for Layer 3 VPNs on page 225](#)

## Understanding VPN-IPv4 Addresses and Route Distinguishers

Because Layer 3 VPNs connect private networks—which can use either public addresses or private addresses, as defined in RFC 1918 (*Address Allocation for Private Internets*)—over the public Internet infrastructure, when the private networks use private addresses, the addresses might overlap with the addresses of another private network.

[Figure 22 on page 218](#) illustrates how private addresses of different private networks can overlap. Here, sites within VPN A and VPN B use the address spaces 10.1.0.0/16, 10.2.0.0/16, and 10.3.0.0/16 for their private networks.

Figure 22: Overlapping Addresses Among Different VPNs



g017158

To avoid overlapping private addresses, you can configure the network devices to use public addresses instead of private addresses. However, this is a large and complex undertaking. The solution provided in RFC 4364 uses the existing private network numbers to create a new address that is unambiguous. The new address is part of the VPN-IPv4 address family, which is a BGP address family added as an extension to the BGP protocol. In VPN-IPv4 addresses, a value that identifies the VPN, called a route distinguisher, is prefixed to the private IPv4 address, providing an address that uniquely identifies a private IPv4 address.

Only the PE routers need to support the VPN-IPv4 address extension to BGP. When an ingress PE router receives an IPv4 route from a device within a VPN, it converts it into a VPN-IPv4 route by adding the route distinguisher prefix to the route. The VPN-IPv4 addresses are used only for routes exchanged between PE routers. When an egress PE router receives a VPN-IPv4 route, it converts the VPN-IPv4 route back to an IPv4 route by removing the route distinguisher before announcing the route to its connected CE routers.

VPN-IPv4 addresses have the following format:

- Route distinguisher is a 6-byte value that you can specify in one of the following formats:
  - **as-number:number**, where **as-number** is an AS number (a 2-byte value) and **number** is any 4-byte value. The AS number can be in the range 1 through 65,535. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned,

nonprivate AS number, preferably the Internet service provider's (ISP's) own or the customer's own AS number.

- ***ip-address:number***, where ***ip-address*** is an IP address (a 4-byte value) and ***number*** is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a nonprivate address in your assigned prefix range.
- IPv4 address—4-byte address of a device within the VPN.

Figure 22 on page 218 illustrates how the AS number can be used in the route distinguisher. Suppose that VPN A is in AS 65535 and that VPN B is in AS 666 (both these AS numbers belong to the ISP), and suppose that the route distinguisher for Site 2 in VPN A is 65535:02 and that the route distinguisher for Site 2 in VPN B is 666:02. When Router PE2 receives a route from the CE router in VPN A, it converts it from its IP address of 10.2.0.0 to a VPN-IPv4 address of 65535:02:10.2.0.0. When the PE router receives a route from VPN B, which uses the same address space as VPN A, it converts it to a VPN-IPv4 address of 666:02:10.2.0.0.

If the IP address is used in the route distinguisher, suppose Router PE2's IP address is 172.168.0.1. When the PE router receives a route from VPN A, it converts it to a VPN-IPv4 address of 172.168.0.1:0:10.2.0.0/16, and it converts a route from VPN B to 172.168.0.0:1:10.2.0.0/16.

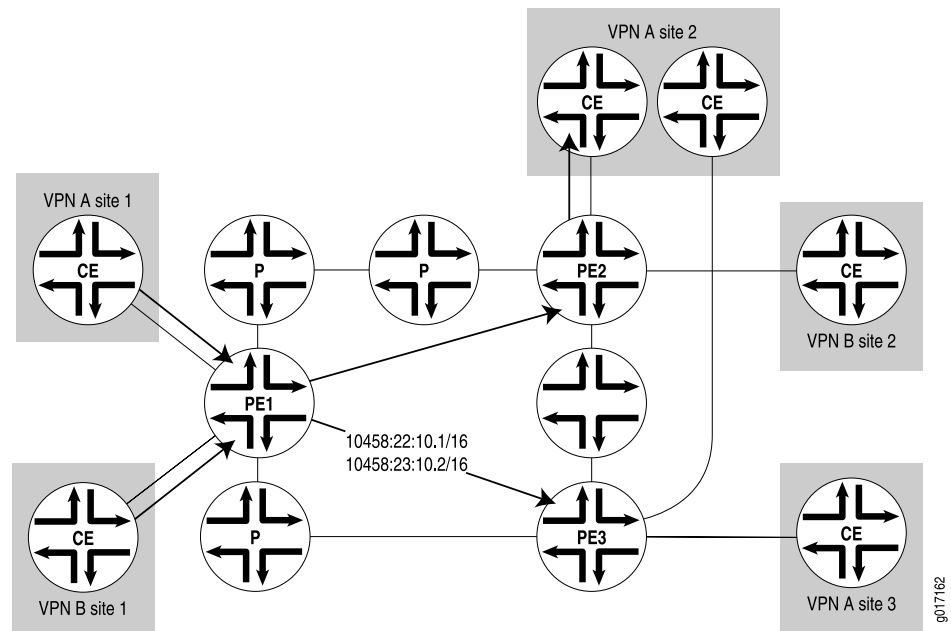
Route distinguishers are used only among PE routers to IPv4 addresses from different VPNs. The ingress PE router creates a route distinguisher and converts IPv4 routes received from CE routers into VPN-IPv4 addresses. The egress PE routers convert VPN-IPv4 routes into IPv4 routes before announcing them to the CE router.

Because VPN-IPv4 addresses are a type of BGP address, you must configure IBGP sessions between pairs of PE routers so that the PE routers can distribute VPN-IPv4 routes within the provider's core network. (All PE routers are assumed to be within the same AS.)

You define BGP communities to constrain the distribution of routes among the PE routers. Defining BGP communities does not, by itself, distinguish IPv4 addresses.

Figure 23 on page 220 illustrates how Router PE1 adds the route distinguisher 10458:22:10.1/16 to routes received from the CE router at Site 1 in VPN A and forwards these routes to the other two PE routers. Similarly, Router PE1 adds the route distinguisher 10458:23:10.2/16 to routes received by the CE router at Site 1 in VPN B and forwards these routes to the other PE routers.

Figure 23: Route Distinguishers

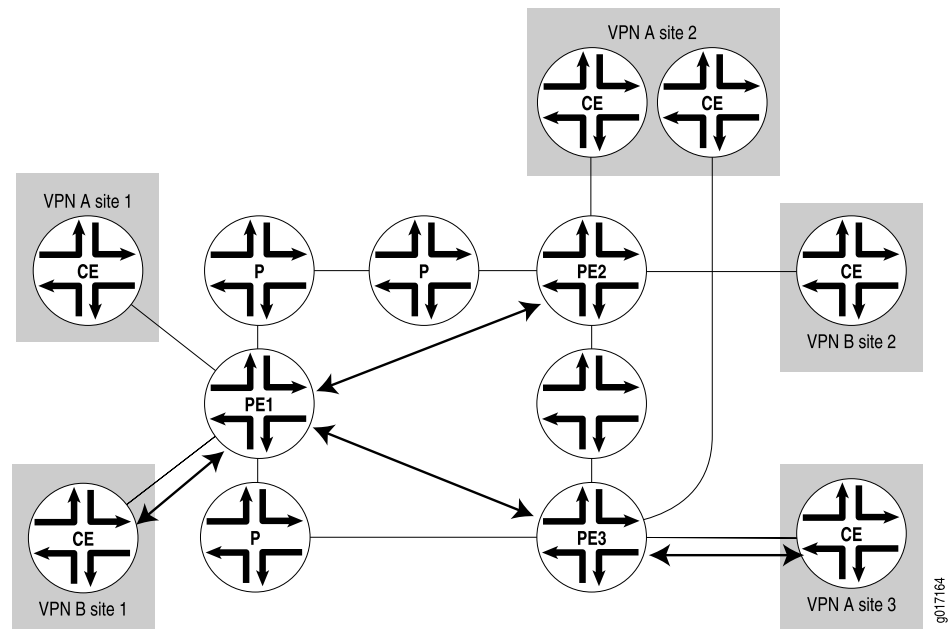




## Understanding IPv4 Route Distribution in a Layer 3 VPN

Within a VPN, the distribution of VPN-IPv4 routes occurs between the PE and CE routers and between the PE routers (see [Figure 24 on page 221](#)).

*Figure 24: Route Distribution Within a VPN*



This section discusses the following topics:

- [Distribution of Routes from CE to PE Routers on page 221](#)
- [Distribution of Routes Between PE Routers on page 222](#)
- [Distribution of Routes from PE to CE Routers on page 224](#)

### Distribution of Routes from CE to PE Routers

A CE router announces its routes to the directly connected PE router. The announced routes are in IPv4 format. The PE router places the routes into the VRF table for the VPN. In the Junos OS, this is the *routing-instance-name.inet.0* routing table, where *routing-instance-name* is the configured name of the VPN.

The connection between the CE and PE routers can be a remote connection (a WAN connection) or a direct connection (such as a Frame Relay or Ethernet connection).

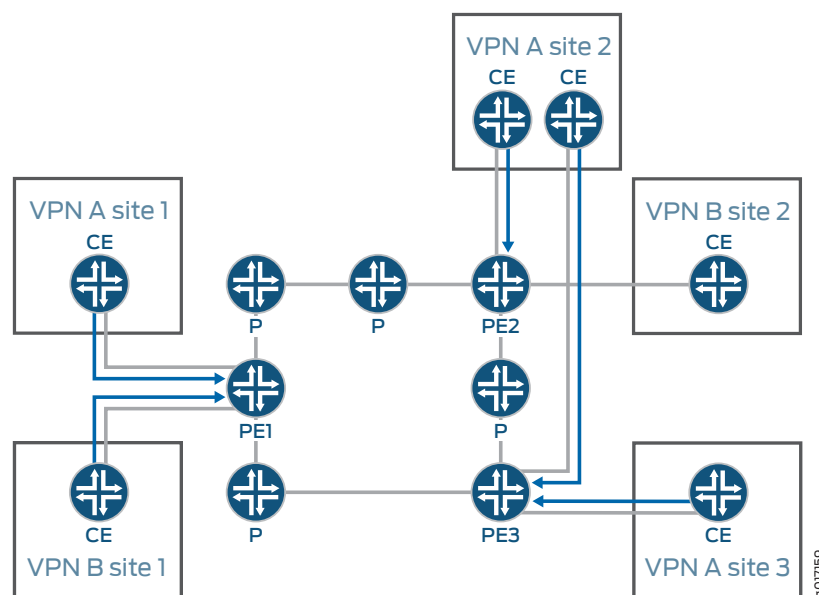
CE routers can communicate with PE routers using one of the following:

- OSPF
- RIP

- BGP
- Static route

Figure 25 on page 222 illustrates how routes are distributed from CE routers to PE routers. Router PE1 is connected to two CE routers that are in different VPNs. Therefore, it creates two VRF tables, one for each VPN. The CE routers announce IPv4 routes. The PE router installs these routes into two different VRF tables, one for each VPN. Similarly, Router PE2 creates two VRF tables into which routes are installed from the two directly connected CE routers. Router PE3 creates one VRF table because it is directly connected to only one VPN.

Figure 25: Distribution of Routes from CE Routers to PE Routers




## Distribution of Routes Between PE Routers

When one PE router receives routes advertised from a directly connected CE router, it checks the received route against the VRF export policy for that VPN. If it matches, the route is converted to VPN-IPv4 format—that is, the 8-byte route distinguisher is prepended to the 4-byte VPN prefix to form a 12-byte VPN-IPv4 address. The route is then tagged with a route target community. The PE router announces the route in VPN-IPv4 format to the remote PE routers for use by VRF import policies. The routes are distributed using IBGP sessions, which are configured in the provider's core network. If the route does not match, it is not exported to other PE routers, but can still be used locally for routing, for example, if two CE routers in the same VPN are directly connected to the same PE router.

The remote PE router places the route into its `bgp.l3vpn.0` table if the route passes the import policy on the IBGP session between the PE routers. At the same time, it checks the route against the VRF import policy for the VPN. If it matches, the route distinguisher is removed from the route, and it is placed into the VRF table (the `routing-instance-name.inet.0` table) in IPv4 format.

Figure 26 on page 223 illustrates how Router PE1 distributes routes to the other PE routers in the provider's core network. Router PE2 and Router PE3 each have VRF import policies that they use to determine whether to accept routes received over the IBGP sessions and install them in their VRF tables.

Figure 26: Distribution of Routes Between PE Routers



When a PE router receives routes advertised from a directly connected CE router (Router PE1 in Figure 26 on page 223), it uses the following procedure to examine the route, convert it to a VPN route, and distribute it to the remote PE routers:

1. The PE router checks the received route using the VRF export policy for that VPN.
2. If the received route matches the export policy, the route is processed as follows:
  - a. The route is converted to VPN-IPv4 format—that is, the 8-byte route distinguisher is prepended to the 4-byte VPN prefix to form the 12-byte VPN-IPv4 address.
  - b. A route target community is added to the route.
  - c. The PE router advertises the route in VPN-IPv4 format to the remote PE routers. The routes are distributed using IBGP sessions, which are configured in the provider's core network.
3. If the route does not match the export policy, it is not exported to the remote PE routers, but can still be used locally for routing—for example, if two CE routers in the same VPN are directly connected to the same PE router.

When the remote PE router receives routes advertised from another PE router (Routers PE2 and PE3 in [Figure 26 on page 223](#)), it uses the following procedure to process the route:

1. If the route is accepted by the import policy on the IBGP session between the PE routers, the remote PE router places the route into its `bgp.l3vpn.0` table.
2. The remote PE router checks the route's route target community against the VRF import policy for the VPN.
3. If it matches, the route distinguisher is removed from the route, and it is placed into the VRF table (the *routing-instance-name*.inet.0 table) in IPv4 format.

### Distribution of Routes from PE to CE Routers


The remote PE router announces the routes in its VRF tables, which are in IPv4 format, to its directly connected CE routers.

PE routers can communicate with CE routers using one of the following routing protocols:

- OSPF
- RIP
- BGP
- Static route

[Figure 27 on page 225](#) illustrates how the three PE routers announce their routes to their connected CE routers.

Figure 27: Distribution of Routes from PE Routers to CE Routers



## Configuring IPv4 Packet Forwarding for Layer 3 VPNs

You can configure the router to support packet forwarding for IPv4 traffic in Layer 2 and Layer 3 VPNs. Packet forwarding is handled in one of the following ways, depending on the type of helper service configured:

- **BOOTP service**—Clients send Bootstrap Protocol (BOOTP) requests through the router configured with BOOTP service to a server in the specified routing instance. The server recognizes the client address and sends a response back to the router configured with BOOTP service. This router forwards the reply to the correct client address in the specified routing instance.
- **Other services**—Clients send requests through the router configured with the service to a server in the specified routing instance. The server recognizes the client address and sends a response to the correct client address in the specified routing instance.

To enable packet forwarding for VPNs, include the **helpers** statement:

```
helpers {
  service {
    description description-of-service;
    server {
      address address {
        routing-instance routing-instance-names;
      }
    }
  }
  interface interface-name {
```

```

description description-of-interface;
no-listen;
server {
    address address {
        routing-instance routing-instance-names;
    }
}
}
}
}

```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options]
- [edit logical-systems *logical-system-name* forwarding-options]
- [edit routing-instances *routing-instance-name* forwarding-options]



**NOTE:** You can enable packet forwarding for multiple VPNs. However, the client and server must be within the same VPN. Any Juniper Networks routing platforms with packet forwarding enabled along the path between the client and server must also reside within the same VPN.

The address and routing instance together constitute a unique server. This has implications for routers configured with BOOTP service, which can accept multiple servers.

For example, a BOOTP service can be configured as follows:

```

[edit forwarding-options helpers bootp]
server address 10.2.3.4 routing-instance [instance-A instance-B];

```

Even though the addresses are identical, the routing instances are different. A packet coming in for BOOTP service on **instance-A** is forwarded to **10.2.3.4** in the **instance-A** routing instance, while a packet coming in on **instance-B** is forwarded in the **instance-B** routing instance. Other services can only accept a single server, so this configuration does not apply in those cases.

#### Related Documentation

- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*

# Carrying IPv6 Traffic Over Layer 3 VPNs

- [Understanding IPv6 Layer 3 VPNs on page 227](#)
- [Configuring Layer 3 VPNs to Carry IPv6 Traffic on page 227](#)
- [Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains on page 231](#)

## Understanding IPv6 Layer 3 VPNs

---

The interfaces between the PE and CE routers of a Layer 3 VPN can be configured to carry IP version 6 (IPv6) traffic. IP allows numerous nodes on different networks to interoperate seamlessly. IPv4 is currently used in intranets and private networks, as well as the Internet. IPv6 is the successor to IPv4, and is based for the most part on IPv4.

In the Juniper Networks implementation of IPv6, the service provider implements an MPLS-enabled IPv4 backbone to provide VPN service for IPv6 customers. The PE routers have both IPv4 and IPv6 capabilities. They maintain IPv6 VPN routing and forwarding (VRF) tables for their IPv6 sites and encapsulate IPv6 traffic in MPLS frames that are then sent into the MPLS core network.

IPv6 for Layer 3 VPNs is supported for BGP and for static routes.

IPv6 over Layer 3 VPNs is described in RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*.

**Related  
Documentation**

- [Routing Policies, Firewall Filters, and Traffic Policers Feature Guide](#)
- [Junos OS Routing Protocols Library](#)

## Configuring Layer 3 VPNs to Carry IPv6 Traffic

---

You can configure IP version 6 (IPv6) between the PE and CE routers of a Layer 3 VPN. The PE router must have the PE router to PE router BGP session configured with the **family inet6-vpn** statement. The CE router must be capable of receiving IPv6 traffic. You can configure BGP or static routes between the PE and CE routers.

The following sections explain how to configure IPv6 VPNs between the PE routers:

- [Configuring IPv6 on the PE Router on page 228](#)
- [Configuring the Connection Between the PE and CE Routers on page 228](#)
- [Configuring IPv6 on the Interfaces on page 230](#)

## Configuring IPv6 on the PE Router

To configure IPv6 between the PE and CE routers, include the **family inet6-vpn** statement in the configuration on the PE router:

```
family inet6-vpn {  
  (any | multicast | unicast) {  
    aggregate-label community community-name;  
    prefix-limit maximum prefix-limit;  
    rib-group rib-group-name;  
  }  
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You also must include the **ipv6-tunneling** statement:

```
ipv6-tunneling;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

## Configuring the Connection Between the PE and CE Routers

To support IPv6 routes, you must configure BGP, OSPF version 3, IS-IS, or static routes for the connection between the PE and CE routers in the Layer 3 VPN. You can configure BGP to handle just IPv6 routes or both IP version 4 (IPv4) and IPv6 routes.

For more information about IPv6, see *Junos OS Routing Protocols Library*.

For more information about IS-IS see *Example: Configuring IS-IS*,

The following sections explain how to configure BGP and static routes:

- [Configuring BGP on the PE Router to Handle IPv6 Routes on page 228](#)
- [Configuring BGP on the PE Router for IPv4 and IPv6 Routes on page 229](#)
- [Configuring OSPF Version 3 on the PE Router on page 229](#)
- [Configuring Static Routes on the PE Router on page 230](#)

---

### Configuring BGP on the PE Router to Handle IPv6 Routes

To configure BGP in the Layer 3 VPN routing instance to handle IPv6 routes, include the **bgp** statement:



```

bgp {
  group group-name {
    local-address IPv6-address;
    family inet6 {
      unicast;
    }
    peer-as as-number;
    neighbor IPv6-address;
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

### Configuring BGP on the PE Router for IPv4 and IPv6 Routes

To configure BGP in the Layer 3 VPN routing instance to handle both IPv4 and IPv6 routes, include the **bgp** statement:

```

bgp {
  group group-name {
    local-address IPv4-address;
    family inet {
      unicast;
    }
    family inet6 {
      unicast;
    }
    peer-as as-number;
    neighbor address;
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

### Configuring OSPF Version 3 on the PE Router

To configure OSPF version 3 in the Layer 3 VPN routing instance to handle IPv6 routes, include the **ospf3** statement:

```

ospf3 {
  area area-id {

```

```

    interface interface-name;
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

For complete configuration guidelines for this statement, see the *Junos OS Routing Protocols Library*.

### Configuring Static Routes on the PE Router

To configure a static route to the CE router in the Layer 3 VPN routing instance, include the **routing-options** statement:

```

routing-options {
  rib routing-table.inet6.0 {
    static {
      defaults {
        static-options;
      }
    }
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

### Configuring IPv6 on the Interfaces

You need to configure IPv6 on the PE router interfaces to the CE routers and on the CE router interfaces to the PE routers.

To configure the interface to handle IPv6 routes, include the **family inet6** statement:

```

family inet6 {
  address ipv6-address;
}

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number*]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

If you have configured the Layer 3 VPN to handle both IPv4 and IPv6 routes, configure the interface to handle both IPv4 and IPv6 routes by including the **unit** statement:

```
unit unit-number {
  family inet {
    address ipv4-address;
  }
  family inet6 {
    address ipv6-address;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

## Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains

This example shows how to configure Junos OS to tunnel IPv6 over a Layer 3 VPN IPv4 network. Internal BGP (IBGP) is used between the customer edge (CE) and provider edge (PE) devices, as described in Internet draft draft-marques-ppvpn-ibgp-version.txt, *RFC2547bis networks using internal BGP as PE-CE protocol*, instead of the more typical external BGP (EBGP) PE-CE connections.

- [Requirements on page 231](#)
- [Overview on page 232](#)
- [Configuration on page 233](#)
- [Verification on page 238](#)

### Requirements

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

All PE routers participating in a Layer 3 VPN with the **independent-domain** statement in its configuration must be running Junos OS Release 6.3 or later.

## Overview

This example shows one method of enabling a router to participate in a customer VPN autonomous-system (AS) domain and to transparently exchange routing information through a Layer 3 VPN without the customer network attributes being visible to the carrier network, and without the carrier network attributes being visible to the customer network.

As an added requirement, the customer network in this example is based on IPv6, while the provider network uses IPv4.

The **independent-domain** feature is useful when customer route attributes need to be transparently forwarded across the VPN network without even the service-provider (SP) AS path appearing in the routes. In a typical Layer 3 VPN, the route attributes such as the originator ID, cluster list, route metric, and AS path are not transparent from one CE device to another CE device.

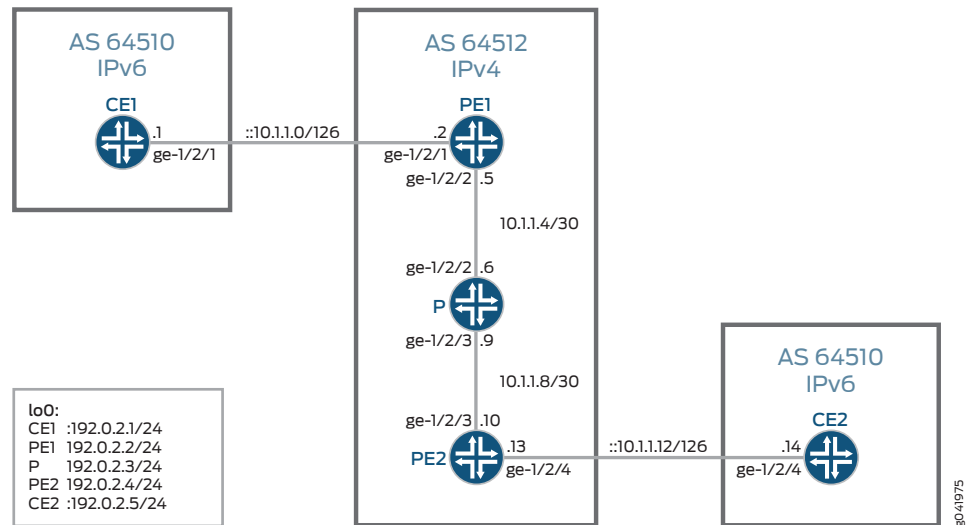
For example, suppose you have a customer VRF whose AS is 1. The customer advertises routes to you through BGP (either IBGP or EBGP). Your core network (the primary routing instance) uses AS 3. Without **independent-domain** configured, if the customer advertises 10.0.0.0/24 to you through BGP, the prefix contains the customer's AS 1 in the AS path. To transport the advertisement across the core to the other PE devices, your core AS 3 is added to the AS path by multiprotocol BGP (MP-BGP). The AS path is now 3 1. When the prefix is advertised out of the core back into the Layer 3 VPN at a remote PE device, the Layer 3 VPN AS 1 is added again, making the AS Path 1 3 1, which is an AS loop. The **independent-domain** statement ensures that only the ASs in the routing-instance are checked during loop detection, and the main, primary routing instances (your core's AS 3) is not considered. This is done by using the attribute 128 (attribute set), which is an optional transitive attribute. The attribute set hides the route's AS path, local preference, and so on, so that those do not appear during the loop check.



**NOTE:** In Junos OS 10.4 and later, you can specify the **no-attrset** option of **independent-domain** so that instead of using attribute 128 (attribute set), Junos OS simply does loop checking on routing-instance ASs without considering your core's AS used in MP-BGP. This is useful if you are using the **local-as** feature, and you only want to configure independent domains to maintain the independence of local ASs in the routing instance, and perform BGP loop detection only for the specified local ASs in the routing instance. In this case, you can disable the attribute set message.

Figure 28 on page 233 shows the sample network.

Figure 28: Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains



“CLI Quick Configuration” on page 233 shows the configuration for all of the devices in Figure 28 on page 233.

The section “Configuring Device PE1” on page 235 describes the steps on Device PE1.

## 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 CE1**

```

set interfaces ge-1/2/1 unit 0 family inet6 address ::10.1.1/126
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet6 address ::192.0.2.1/24
set protocols bgp group toPE1 type internal
set protocols bgp group toPE1 family inet6 unicast
set protocols bgp group toPE1 export send-direct
set protocols bgp group toPE1 neighbor ::10.1.1.2
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
set routing-options router-id 192.0.2.1
set routing-options autonomous-system 64510

```

**Device CE2**

```

set interfaces ge-1/2/4 unit 0 family inet6 address ::10.1.1.14/126
set interfaces ge-1/2/4 unit 0 family mpls
set interfaces lo0 unit 0 family inet6 address ::192.0.2.5/24
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 family inet6 unicast
set protocols bgp group toPE2 export send-direct
set protocols bgp group toPE2 neighbor ::10.1.1.13
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept

```

```
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 64510
```

```
Device PE1  set interfaces ge-1/2/1 unit 0 family inet6 address ::10.1.1.2/126
set interfaces ge-1/2/2 unit 0 family inet address 10.1.1.5/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.2/24
set protocols mpls ipv6-tunneling
set protocols mpls interface ge-1/2/2.0
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 192.0.2.2
set protocols bgp group toPE2 family inet6-vpn unicast
set protocols bgp group toPE2 neighbor 192.0.2.4
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/2.0
set protocols ldp interface ge-1/2/2.0
set protocols ldp interface lo0.0
set routing-instances red instance-type vrf
set routing-instances red interface ge-1/2/1.0
set routing-instances red route-distinguisher 64512:1
set routing-instances red vrf-target target:64512:1
set routing-instances red routing-options router-id 192.0.2.2
set routing-instances red routing-options autonomous-system 64510
set routing-instances red routing-options autonomous-system independent-domain
set routing-instances red protocols bgp group toCE1 type internal
set routing-instances red protocols bgp group toCE1 family inet6 unicast
set routing-instances red protocols bgp group toCE1 neighbor ::10.1.1.1
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 64512
```

```
Device P    set interfaces ge-1/2/2 unit 0 family inet address 10.1.1.6/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces ge-1/2/3 unit 0 family inet address 10.1.1.9/30
set interfaces ge-1/2/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.3/24
set protocols mpls interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set protocols ldp interface all
set routing-options router-id 192.0.2.3
```

```
Device PE2  set interfaces ge-1/2/3 unit 0 family inet address 10.1.1.10/30
set interfaces ge-1/2/3 unit 0 family mpls
set interfaces ge-1/2/4 unit 0 family inet6 address ::10.1.1.13/126
set interfaces lo0 unit 0 family inet address 192.0.2.4/24
set protocols mpls ipv6-tunneling
set protocols mpls interface ge-1/2/3.0
set protocols bgp group toPE1 type internal
set protocols bgp group toPE1 local-address 192.0.2.4
set protocols bgp group toPE1 family inet6-vpn unicast
set protocols bgp group toPE1 neighbor 192.0.2.2
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/3.0
set protocols ldp interface ge-1/2/3.0
```

```

set protocols ldp interface lo0.0
set routing-instances red instance-type vrf
set routing-instances red interface ge-1/2/4.0
set routing-instances red route-distinguisher 64512:1
set routing-instances red vrf-target target:64512:1
set routing-instances red routing-options router-id 192.0.2.4
set routing-instances red routing-options autonomous-system 64510
set routing-instances red routing-options autonomous-system independent-domain
set routing-instances red protocols bgp group toCE2 type internal
set routing-instances red protocols bgp group toCE2 family inet6 unicast
set routing-instances red protocols bgp group toCE2 neighbor ::10.1.1.14
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 64512

```

### Configuring Device PE1

**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 Device PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set ge-1/2/1 unit 0 family inet6 address ::10.1.1.2/126

user@PE1# set ge-1/2/2 unit 0 family inet address 10.1.1.5/30
user@PE1# set ge-1/2/2 unit 0 family mpls

user@PE1# set lo0 unit 0 family inet address 192.0.2.2/24

```

2. Configure MPLS on the interfaces.

```

[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set interface ge-1/2/2.0

```

3. Configure BGP.

```

[edit protocols bgp group toPE2]
user@PE1# set type internal
user@PE1# set local-address 192.0.2.2
user@PE1# set family inet6-vpn unicast
user@PE1# set neighbor 192.0.2.4

```

4. Configure an interior gateway protocol (IGP).

```

[edit protocols ospf area 0.0.0.0]
user@PE1# set interface lo0.0 passive
user@PE1# set interface ge-1/2/2.0

```

5. Configure a signaling protocol.

```
[edit protocols]
user@PE1# set ldp interface ge-1/2/2.0
user@PE1# set ldp interface lo0.0
```

6. Configure the routing instance.

```
[edit routing-instances red]
user@PE1# set instance-type vrf
user@PE1# set interface ge-1/2/1.0
user@PE1# set route-distinguisher 64512:1
user@PE1# set vrf-target target:64512:1
user@PE1# set routing-options router-id 192.0.2.2
user@PE1# set protocols bgp group toCE1 type internal
user@PE1# set protocols bgp group toCE1 family inet6 unicast
user@PE1# set protocols bgp group toCE1 neighbor ::10.1.1.1
```

7. In the routing instance, include the AS number of the customer network, and include the **independent-domain** statement.

```
[edit routing-instances red routing-options]
user@PE1# set autonomous-system 64510
user@PE1# set autonomous-system independent-domain
```

8. In the main instance, configure the router ID and the provider AS number.

```
[edit routing-options]
user@PE1# set router-id 192.0.2.2
user@PE1# set autonomous-system 64512
```

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

```
user@PE1# show interfaces
interfaces {
  ge-1/2/1 {
    unit 0 {
      family inet6 {
        address ::10.1.1.2/126;
      }
    }
  }
  ge-1/2/2 {
    unit 0 {
      family inet {
        address 10.1.1.5/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
```



```
        family inet {
            address 192.0.2.2/24;
        }
    }
}

user@PE1# show protocols
mpls {
    ipv6-tunneling;
    interface ge-1/2/2.0;
}
bgp {
    group toPE2 {
        type internal;
        local-address 192.0.2.2;
        family inet6-vpn {
            unicast;
        }
        neighbor 192.0.2.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface ge-1/2/2.0;
    }
}
ldp {
    interface ge-1/2/2.0;
    interface lo0.0;
}

user@PE1# show routing-instances
red {
    instance-type vrf;
    interface ge-1/2/1.0;
    route-distinguisher 64512:1;
    vrf-target target:64512:1;
    routing-options {
        router-id 192.0.2.2;
        autonomous-system 64510 independent-domain;
    }
    protocols {
        bgp {
            group toCE1 {
                type internal;
                family inet6 {
                    unicast;
                }
                neighbor ::10.1.1.1;
            }
        }
    }
}
```

```

user@PE1# show routing-options
router-id 192.0.2.2;
autonomous-system 64512;

```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying That the CE Devices Have Connectivity on page 238](#)
- [Checking the AS Paths on page 238](#)

### Verifying That the CE Devices Have Connectivity

**Purpose** Make sure that the tunnel is operating.

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

```

user@CE1> ping ::192.0.2.5

PING6(56=40+8+8 bytes) ::10.1.1.1 --> ::192.0.2.5
16 bytes from ::192.0.2.5, icmp_seq=0 hlim=63 time=1.943 ms
16 bytes from ::192.0.2.5, icmp_seq=1 hlim=63 time=1.587 ms
^C
--- ::192.0.2.5 ping6 statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.587/1.765/1.943/0.178 ms

user@CE2> ping ::192.0.2.1

PING6(56=40+8+8 bytes) ::10.1.1.14 --> ::192.0.2.1
16 bytes from ::192.0.2.1, icmp_seq=0 hlim=63 time=2.097 ms
16 bytes from ::192.0.2.1, icmp_seq=1 hlim=63 time=1.610 ms
^C
--- ::192.0.2.1 ping6 statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.610/1.853/2.097/0.244 ms

```

**Meaning** The IPv6 CE devices can communicate over the core IPv4 network.

### Checking the AS Paths

**Purpose** Make sure that the provider AS number does not appear in the CE device routing tables.

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

```

user@CE1> show route protocol bgp detail

inet6.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
::192.0.2.5/24 (1 entry, 1 announced)
    *BGP      Preference: 170/-101

```

```

Next hop type: Indirect
Address: 0x9514354
Next-hop reference count: 6
Source: ::10.1.1.2
Next hop type: Router, Next hop index: 924
Next hop: ::10.1.1.2 via ge-1/2/1.0, selected
Session Id: 0x500001
Protocol next hop: ::10.1.1.2
Indirect next hop: 0x971c000 262147 INH Session ID: 0x500002
State: <Active Int Ext>
Local AS: 64510 Peer AS: 64510
Age: 50:58      Metric2: 0
Validation State: unverified
Task: BGP_64510::10.1.1.2+45824
Announcement bits (2): 0-KRT 2-Resolve tree 2
AS path: I
Accepted
Localpref: 100
Router ID: 192.0.2.2

::10.1.1.12/126 (1 entry, 1 announced)
*BGP      Preference: 170/-101
Next hop type: Indirect
Address: 0x9514354
Next-hop reference count: 6
Source: ::10.1.1.2
Next hop type: Router, Next hop index: 924
Next hop: ::10.1.1.2 via ge-1/2/1.0, selected
Session Id: 0x500001
Protocol next hop: ::10.1.1.2
Indirect next hop: 0x971c000 262147 INH Session ID: 0x500002
State: <Active Int Ext>
Local AS: 64510 Peer AS: 64510
Age: 50:58      Metric2: 0
Validation State: unverified
Task: BGP_64510::10.1.1.2+45824
Announcement bits (2): 0-KRT 2-Resolve tree 2
AS path: I
Accepted
Localpref: 100
Router ID: 192.0.2.2

```

user@CE2> show route protocol bgp detail

```

inet6.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
::192.0.2.1/24 (1 entry, 1 announced)
*BGP      Preference: 170/-101
Next hop type: Indirect
Address: 0x9514354
Next-hop reference count: 6
Source: ::10.1.1.13
Next hop type: Router, Next hop index: 914
Next hop: ::10.1.1.13 via ge-1/2/4.0, selected
Session Id: 0x400001
Protocol next hop: ::10.1.1.13
Indirect next hop: 0x971c000 262150 INH Session ID: 0x400002
State: <Active Int Ext>
Local AS: 64510 Peer AS: 64510
Age: 50:41      Metric2: 0

```

```
Validation State: unverified
Task: BGP_64510::10.1.1.13+59329
Announcement bits (2): 0-KRT 2-Resolve tree 2
AS path: I
Accepted
Localpref: 100
Router ID: 192.0.2.4

::10.1.1.0/126 (1 entry, 1 announced)
*BGP Preference: 170/-101
Next hop type: Indirect
Address: 0x9514354
Next-hop reference count: 6
Source: ::10.1.1.13
Next hop type: Router, Next hop index: 914
Next hop: ::10.1.1.13 via ge-1/2/4.0, selected
Session Id: 0x400001
Protocol next hop: ::10.1.1.13
Indirect next hop: 0x971c000 262150 INH Session ID: 0x400002
State: <Active Int Ext>
Local AS: 64510 Peer AS: 64510
Age: 50:41 Metric2: 0
Validation State: unverified
Task: BGP_64510::10.1.1.13+59329
Announcement bits (2): 0-KRT 2-Resolve tree 2
AS path: I
Accepted
Localpref: 100
Router ID: 192.0.2.4
```

**Meaning** The output shows that for the BGP routes on the CE devices, the AS path attribute does not include the provider AS 64512.

**Related Documentation**

- *Configuring the Ingress Router for MPLS-Signaled LSPs*
- *Minimum RSVP Configuration*

## Configuring an AS for Layer 3 VPNs

- [Configuring Layer 3 VPNs to Carry IBGP Traffic on page 241](#)
- [Example: Configuring a Layer 3 VPN with Route Reflection and AS Override on page 243](#)
- [Configuring the Algorithm That Determines the Active Route to Evaluate AS Numbers in AS Paths for VPN Routes on page 253](#)

### Configuring Layer 3 VPNs to Carry IBGP Traffic

An independent AS domain is separate from the primary routing instance domain. An AS is a set of routers that are under a single technical administration and that generally use a single IGP and metrics to propagate routing information within the set of routers. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

Configuring an independent domain allows you to keep the AS paths of the independent domain from being shared with the AS path and AS path attributes of other domains, including the master routing instance domain.

If you are using BGP on the router, you must configure an AS number.

When you configure BGP as the routing protocol between a PE router and a CE router in a Layer 3 VPN, you typically configure external peering sessions between the Layer 3 VPN service provider and the customer network ASs.

If the customer network has several sites advertising routes through an external BGP session to the service provider network and if the same AS is used by all the customer sites, the CE routers reject routes from the other CE routers. They detect a loop in the BGP AS path attribute.

To prevent the CE routers from rejecting each other's routes, you could configure the following:

- PE routers advertising routes received from remote PE routers can remap the customer network AS number to its own AS number.
- AS path loops can be configured.
- The customer network can be configured with different AS numbers at each site.

These types of configurations can work when there are no BGP routing exchanges between the customer network and other networks. However, they do have limitations for customer networks that use BGP internally for purposes other than carrying traffic between the CE routers and the PE routers. When those routes are advertised outside the customer network, the service provider ASs are present in the AS path.

To improve the transparency of Layer 3 VPN services for customer networks, you can configure the routing instance for the Layer 3 VPN to isolate the customer's network attributes from the service provider's network attributes.

When you include the **independent-domain** statement in the Layer 3 VPN routing instance configuration, BGP attributes received from the customer network (from the CE router) are stored in a BGP attribute (ATTRSET) that functions like a stack. When that route is advertised from the remote PE router to the remote CE router, the original BGP attributes are restored. This is the default behavior for BGP routes that are advertised to Layer 3 VPNs located in different domains.

This functionality is described in the Internet draft *draft-marques-ppvpn-ibgp-version.txt*, *RFC 2547bis Networks Using Internal BGP as PE-CE Protocol*.

To allow a Layer 3 VPN to transport IBGP traffic, include the **independent-domain** statement:

**independent-domain;**

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* routing-options autonomous-system *number*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options autonomous-system *number*]



**NOTE:** All PE routers participating in a Layer 3 VPN with the **independent-domain** statement in its configuration must be running Junos OS Release 6.3 or later.



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

The independent domain uses the transitive path attribute 128 (attribute set) to tunnel the independent domain's BGP attributes through the Internal BGP (IBGP) core. In Junos OS Release 10.3 and later, if BGP receives attribute 128 and you have not configured an independent domain in any routing instance, BGP treats the received attribute 128 as an unknown attribute.

There is a limit of 16 ASs for each domain.

- Related Documentation**
- [Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains on page 231](#)
  - [Disabling Attribute Set Messages on Independent AS Domains for BGP Loop Detection](#)

## Example: Configuring a Layer 3 VPN with Route Reflection and AS Override

Suppose that you are a service provider providing a managed MPLS-based Layer 3 VPN service. Your customer has several sites and requires BGP routing to customer edge (CE) devices at each site.

- [Requirements on page 243](#)
- [Overview on page 243](#)
- [Configuration on page 244](#)
- [Verification on page 251](#)

### Requirements

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


### Overview

This example has two CE devices, two provider edge (PE) devices, and several provider core devices. The provider network is also using IS-IS to support LDP and BGP loopback reachability. Device P2 is acting as a route reflector (RR). Both CE devices are in autonomous system (AS) 64512. The provider network is in AS 65534.

The **as-override** statement is applied to the PE devices, thus replacing the CE device's AS number with that of the PE device. This prevents the customer AS number from appearing more than once in the AS path attribute.

Figure 29 on page 243 shows the topology used in this example.

Figure 29: AS Override Topology



[“CLI Quick Configuration” on page 244](#) shows the configuration for all of the devices in [Figure 29 on page 243](#). The section [“Step-by-Step Procedure” on page 247](#) describes the steps on Device PE1.

## 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 CE1**

```
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 10.255.1.1/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0101.00
set protocols bgp group PE type external
set protocols bgp group PE family inet unicast
set protocols bgp group PE export ToBGP
set protocols bgp group PE peer-as 65534
set protocols bgp group PE neighbor 10.0.0.2
set policy-options policy-statement ToBGP term Direct from protocol direct
set policy-options policy-statement ToBGP term Direct then accept
set routing-options router-id 10.255.1.1
set routing-options autonomous-system 64512
```

**Device P1**

```
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 family inet address 10.0.0.25/30
set interfaces ge-1/2/2 unit 0 family iso
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.3.3/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0303.00
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.3.3
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.4.4
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-options router-id 10.255.3.3
```



Device P2

```

set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0404.00
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group Core-RRClients type internal
set protocols bgp group Core-RRClients local-address 10.255.4.4
set protocols bgp group Core-RRClients family inet-vpn unicast
set protocols bgp group Core-RRClients cluster 10.255.4.4
set protocols bgp group Core-RRClients peer-as 65534
set protocols bgp group Core-RRClients neighbor 10.255.3.3
set protocols bgp group Core-RRClients neighbor 10.255.7.7
set protocols bgp group Core-RRClients neighbor 10.255.2.2
set protocols bgp group Core-RRClients neighbor 10.255.5.5
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-options router-id 10.255.4.4
set routing-options autonomous-system 65534

```

Device P3

```

set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 family inet address 10.0.0.30/30
set interfaces ge-1/2/2 unit 0 family iso
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.7.7/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0707.00
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.7.7
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.4.4
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable

```

```
set routing-options router-id 10.255.7.7
```

```
Device PE1
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.5/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 family inet address 10.0.0.21/30
set interfaces ge-1/2/2 unit 0 family iso
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.2.2/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0202.00
set protocols mpls interface ge-1/2/2.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.2.2
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.4.4
set protocols isis interface ge-1/2/1.0 level 2 metric 10
set protocols isis interface ge-1/2/1.0 level 1 disable
set protocols isis interface ge-1/2/2.0 level 2 metric 10
set protocols isis interface ge-1/2/2.0 level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface ge-1/2/1.0
set protocols ldp interface ge-1/2/2.0
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface ge-1/2/0.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A protocols bgp group CE type external
set routing-instances VPN-A protocols bgp group CE family inet unicast
set routing-instances VPN-A protocols bgp group CE neighbor 10.0.0.1 peer-as 64512
set routing-instances VPN-A protocols bgp group CE neighbor 10.0.0.1 as-override
set routing-options router-id 10.255.2.2
set routing-options autonomous-system 65534
```

```
Device PE2
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces ge-1/2/2 unit 0 family iso
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.5.5/32
```

```

set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0505.00
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/2.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.5.5
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.4.4
set protocols isis interface ge-1/2/0.0 level 2 metric 10
set protocols isis interface ge-1/2/0.0 level 1 disable
set protocols isis interface ge-1/2/2.0 level 2 metric 10
set protocols isis interface ge-1/2/2.0 level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface ge-1/2/0.0
set protocols ldp interface ge-1/2/2.0
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface ge-1/2/1.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A protocols bgp group CE type external
set routing-instances VPN-A protocols bgp group CE family inet unicast
set routing-instances VPN-A protocols bgp group CE neighbor 10.0.0.18 peer-as 64512
set routing-instances VPN-A protocols bgp group CE neighbor 10.0.0.18 as-override
set routing-options router-id 10.255.5.5
set routing-options autonomous-system 65534

```

<b>Device CE2</b>	<pre> set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.18/30 set interfaces ge-1/2/0 unit 0 family iso set interfaces lo0 unit 0 family inet address 10.255.6.6/32 set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0606.00 set protocols bgp group PE type external set protocols bgp group PE family inet unicast set protocols bgp group PE export ToBGP set protocols bgp group PE peer-as 65534 set protocols bgp group PE neighbor 10.0.0.17 set policy-options policy-statement ToBGP term Direct from protocol direct set policy-options policy-statement ToBGP term Direct then accept set routing-options router-id 10.255.6.6 set routing-options autonomous-system 64512 </pre>
-------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**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 AS override:

1. Configure the interfaces.

To enable MPLS, include the protocol family on the interface so that the interface does not discard incoming MPLS traffic.

```
[edit interfaces]
user@PE1# set ge-1/2/0 unit 0 family inet address 10.0.0.2/30
user@PE1# set ge-1/2/0 unit 0 family iso
user@PE1# set ge-1/2/0 unit 0 family mpls
user@PE1# set ge-1/2/1 unit 0 family inet address 10.0.0.5/30
user@PE1# set ge-1/2/1 unit 0 family iso
user@PE1# set ge-1/2/1 unit 0 family mpls
user@PE1# set ge-1/2/2 unit 0 family inet address 10.0.0.21/30
user@PE1# set ge-1/2/2 unit 0 family iso
user@PE1# set ge-1/2/2 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 10.255.2.2/32
user@PE1# set lo0 unit 0 family iso address 49.0001.0010.0000.0202.00
```

2. Add the interface to the MPLS protocol to establish the control plane level connectivity.

Set up the IGP so that the provider devices can communicate with each other.

To establish a mechanism to distribute MPLS labels, enable LDP. Optionally, for LDP, enable forwarding equivalence class (FEC) deaggregation, which results in faster global convergence.

```
[edit protocols]
user@PE1# set mpls interface ge-1/2/2.0
user@PE1# set mpls interface ge-1/2/1.0
user@PE1# set mpls interface lo0.0
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set isis interface ge-1/2/1.0 level 2 metric 10
user@PE1# set isis interface ge-1/2/1.0 level 1 disable
user@PE1# set isis interface ge-1/2/2.0 level 2 metric 10
user@PE1# set isis interface ge-1/2/2.0 level 1 disable
user@PE1# set isis interface fxp0.0 disable
user@PE1# set isis interface lo0.0 level 2 metric 0
user@PE1# set ldp deaggregate
user@PE1# set ldp interface ge-1/2/1.0
user@PE1# set ldp interface ge-1/2/2.0
user@PE1# set ldp interface fxp0.0 disable
user@PE1# set ldp interface lo0.0
```

3. Enable the internal BGP (IBGP) connection to peer with the RR only, using the IPv4 VPN unicast address family.

```
[edit protocols bgp group l3vpn]
user@PE1# set type internal
user@PE1# set local-address 10.255.2.2
user@PE1# set family inet-vpn unicast
user@PE1# set peer-as 65534
user@PE1# set local-as 65534
user@PE1# set neighbor 10.255.4.4
```

4. Configure the routing instance, including the **as-override** statement.

Create the routing-instance (VRF) on the PE device, setting up the BGP configuration to peer with Device CE1.

```
[edit routing-instances VPN-A]
user@PE1# set instance-type vrf
user@PE1# set interface ge-1/2/0.0
user@PE1# set route-distinguisher 65534:1234
user@PE1# set vrf-target target:65534:1234
user@PE1# set protocols bgp group CE type external
user@PE1# set protocols bgp group CE family inet unicast
user@PE1# set protocols bgp group CE neighbor 10.0.0.1 peer-as 64512
user@PE1# set protocols bgp group CE neighbor 10.0.0.1 as-override
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@PE1# set router-id 10.255.2.2
user@PE1# set autonomous-system 65534
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **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@PE1# show interfaces
ge-1/2/0 {
  unit 2 {
    family inet {
      address 10.0.0.2/30;
    }
    family iso;
    family mpls;
  }
}
ge-1/2/1 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
    family iso;
    family mpls;
  }
}
ge-1/2/2 {
  unit 21 {
    family inet {
      address 10.0.0.21/30;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
```

```
        family inet {
            address 10.255.2.2/32;
        }
        family iso {
            address 49.0001.0010.0000.0202.00;
        }
    }
}

user@PE1# show protocols
mpls {
    interface ge-1/2/2.0;
    interface ge-1/2/1.0;
    interface lo0.0;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group l3vpn {
        type internal;
        local-address 10.255.2.2;
        family inet-vpn {
            unicast;
        }
        peer-as 65534;
        local-as 65534;
        neighbor 10.255.4.4;
    }
}
isis {
    interface ge-1/2/1.0 {
        level 2 metric 10;
        level 1 disable;
    }
    interface ge-1/2/2.0 {
        level 2 metric 10;
        level 1 disable;
    }
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        level 2 metric 0;
    }
}
ldp {
    deaggregate;
    interface ge-1/2/1.0;
    interface ge-1/2/2.0;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}

user@PE1# show routing-instances
```

```

VPN-A {
  instance-type vrf;
  interface ge-1/2/0.0;
  route-distinguisher 65534:1234;
  vrf-target target:65534:1234;
  protocols {
    bgp {
      group CE {
        type external;
        family inet {
          unicast;
        }
        neighbor 10.0.0.1 {
          peer-as 64512;
          as-override;
        }
      }
    }
  }
}

```

```

user@PE1# show routing-options
router-id 10.255.2.2;
autonomous-system 65534;

```

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

## Verification

Confirm that the configuration is working properly.

- [Checking AS Path to the CE Devices on page 251](#)
- [Checking How the Route to Device CE2 Is Advertised on page 252](#)
- [Checking the Route on Device CE1 on page 252](#)

### Checking AS Path to the CE Devices

<b>Purpose</b>	Display information on Device PE1 about the AS path attribute for the route to Device CE2's loopback interface.
<b>Action</b>	On Device PE1, from operational mode, enter the <b>show route table VPN-A.inet.0 10.255.6.6</b> command.
<pre> user@PE1&gt; show route table VPN-A.inet.0 10.255.6.6 </pre>	
<pre> VPN-A.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both </pre>	
<b>10.255.6.6/32</b>	<pre> *[BGP/170] 02:19:35, localpref 100, from 10.255.4.4   AS path: 64512 I, validation-state: unverified   &gt; to 10.0.0.22 via ge-1/2/2.0, Push 300032, Push 299776(top) </pre>

**Meaning** The output shows that Device PE1 has an AS path for 10.255.6.6/32 as coming from AS 64512.

### Checking How the Route to Device CE2 Is Advertised

**Purpose** Make sure the route to Device CE2 is advertised to Device CE1 as if it is coming from the MPLS core.

**Action** On Device PE1, from operational mode, enter the **show route advertising-protocol bgp 10.0.0.1** command.

```
user@PE1> show route advertising-protocol bgp 10.0.0.1
```

```
VPN-A.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lc1pref  AS path
* 10.0.0.16/30          Self
* 10.255.1.1/32         10.0.0.1        65534 I
* 10.255.6.6/32        Self             65534 I
```

**Meaning** The output indicates that Device PE1 is advertising only its own AS number in the AS path.

### Checking the Route on Device CE1

**Purpose** Make sure that Device CE1 contains only the provider AS number in the AS path for the route to Device CE2.

**Action** From operational mode, enter the **show route table inet.0 terse 10.255.6.6** command.

```
user@CE1> show route table inet.0 terse 10.255.6.6
```

```
inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
A V Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* ? 10.255.6.6/32    B 170    100          >10.0.0.2    65534 65534
I
unverified
```

**Meaning** The output indicates that Device CE1 has a route to Device CE2. The loop issue is resolved with the use of the **as-override** statement.

One route is hidden on the CE device. This is because Junos OS does not perform a BGP split horizon. Generally, split horizon in BGP is unnecessary, because any routes that might be received back by the originator are less preferred due to AS path length (for EBGP), AS path loop detection (IBGP), or other BGP metrics. Advertising routes back to the neighbor from which they were learned has a negligible effect on the router's performance, and is the correct thing to do.



- Related Documentation**
- [Understanding AS Override](#)
  - [Example: Enabling BGP Route Advertisements](#)

## Configuring the Algorithm That Determines the Active Route to Evaluate AS Numbers in AS Paths for VPN Routes

By default, the third step of the algorithm that determines the active route evaluates the length of the AS path but not the contents of the AS path. In some VPN scenarios with BGP multiple path routes, it can also be useful to compare the AS numbers of the AS paths and to have the algorithm select the route whose AS numbers match.

To configure the algorithm that selects the active path to evaluate the AS numbers in AS paths for VPN routes:

- Include the **as-path-compare** statement at the **[edit routing-instances routing-instance-name routing-options multipath]** hierarchy level.



**NOTE:** The **as-path-compare** statement is not supported for the default routing instance.

- Related Documentation**
- [as-path-compare on page 1044](#)



## CHAPTER 17

# Limiting Routes Using Route Resolution and Route Filters

- [Example: Configuring Route Resolution on PE Routers on page 255](#)
- [Example: Configuring Route Resolution on Route Reflectors on page 257](#)
- [Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs on page 260](#)

## Example: Configuring Route Resolution on PE Routers

---

This example shows how to configure a routing table to accept routes from specific routing tables. It also shows how to configure a routing table to use specific import policies to produce a route resolution table to resolve routes.

- [Requirements on page 255](#)
- [Overview on page 255](#)
- [Configuration on page 256](#)
- [Verification on page 257](#)

### Requirements

Before you begin, configure a Layer 3 VPN, as shown in one of the following examples:

- [Example: Configuring Interprovider Layer 3 VPN Option A on page 334](#)
- [Example: Configuring Interprovider Layer 3 VPN Option B on page 355](#)

### Overview

One method to achieve IPv4 route scaling is to modify how BGP routes are added to the forwarding tables. By default, the Routing Protocol Process (rpd) adds all the routes in inet.0 and inet.3 to the resolution tree. Normally, this includes the resolved IPv4 BGP routes, which can increase memory consumption. To achieve better scaling for IPv4 routes, this example shows how to configure the Junos OS so that resolved BGP routes are not added to the resolution tree. This is achieved by applying an import policy on the route resolution table, which ensures it does not accept any BGP routes from inet.0.

You would apply this configuration to all provider edge (PE) routers in the Layer 3 VPN.

## 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.

**PE Router**

```
set policy-options policy-statement protocol-bgp from protocol bgp
set policy-options policy-statement protocol-bgp then reject
set routing-options resolution rib inet.0 import protocol-bgp
```

**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 route resolution:

1. Configure the routing policy.  

```
[edit policy-options policy-statement protocol-bgp]
user@PE# set from protocol bgp
user@PE# set then reject
```
2. Apply the routing policy.  

```
[edit routing-options resolution]
user@PE# set rib inet.0 import protocol-bgp
```
3. If you are done configuring the device, commit the configuration.  

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

---

## Results

Confirm your configuration by issuing the **show policy-options** and **show routing-options** commands.

```
user@PE# show policy-options
policy-statement protocol-bgp {
  from protocol bgp;
  then reject;
}

user@PE# show routing-options
resolution {
  rib inet.0 {
    import protocol-bgp;
  }
}
```

## Verification

Confirm that the configuration is working properly by running the following commands:

- *show route*
- *show route forwarding-table*
- *show route resolution*

### Related Documentation

- [Example: Configuring Route Resolution on Route Reflectors on page 257](#)

---

## Example: Configuring Route Resolution on Route Reflectors

This example shows how to change the default resolution behavior on a route reflector (RR) to use inet.0 for next-hop resolution instead of inet.3.

- [Requirements on page 257](#)
- [Overview on page 257](#)
- [Configuration on page 258](#)
- [Verification on page 260](#)

## Requirements

Before you begin, configure a Layer 3 VPN, as shown in one of the following examples:

- [Example: Configuring Interprovider Layer 3 VPN Option A on page 334](#)
- [Example: Configuring Interprovider Layer 3 VPN Option B on page 355](#)

## Overview

One scenario for route resolution is when you have a label-switched path configured from an RR to a provider edge (PE) router, or when the PE routers only peer with the RR. This can result in routes being hidden. To resolve this issue, you can change the default resolution behavior to use inet.0 for next-hop resolution.

By default, the `bgp.l3vpn.0` routing table stores all VPN-IPv4 unicast routes. This table is present on any router that has Layer 3 VPNs configured, including PE routers and RRs.

When a Layer 3 VPN router receives a route from another Layer 3 VPN router, it places the route into its `bgp.l3vpn.0` routing table. The route is resolved using the information in the `inet.3` routing table. This means that when BGP receives a route destined for table `bgp.l3vpn.0`, the protocol nexthop (received BGP nexthop) has its forwarding nexthop recursively determined from the `inet.3` table. The resulting route is converted into IPv4 format and redistributed to all `routing-instance-name.inet.0` routing tables if it matches the VRF import policy.

On an RR with no attached customer edge (CE) routers, the **resolution rib `bgp.l3vpn.0` resolution-ribs `inet.0`** configuration causes routes in `bgp.l3vpn.0` to use the information

in `inet.0` instead of `inet.3` to resolve routes. You should not use this configuration on a router that is directly attached to a CE router. In other words, do not use **`resolution rib bgp.l3vpn.0 resolution-ribs inet.0`** on a PE router.

If you want both `inet.0` and `inet.3` to be used, you must configure both, as in **`set resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3]`**.

In this example, the policy **`POLICY-limit-resolve-routes`** limits the route resolution to only routes learned through IS-IS. If you omit the import policy, all routes in `inet.0` are evaluated and potentially used to resolve the protocol next hop. If you do not want to resolve against all entries, you use a policy to filter for a subset of the routes from the tables that are used for route resolution.

One common example is when you resolve against all routes in `inet.0`, except the default route (0/0).

Although the **`import`** statement is used in this configuration, no routes are imported or copied. Rather, the **`import policy-name`** configuration limits the set of possible routes that can be considered for route resolution.

The **`resolution rib bgp.l3vpn.0 resolution-ribs inet.0`** configuration is useful when a BGP RR is not in the forwarding path. In other words, there are no ingress LSPs at the RR. Consider the case where RSVP is the label signaling protocol, and RSVP is configured full mesh at the edge routers. The RR needs to be able to reflect the routes. To do so, BGP is expected to perform a route resolvability check. If a Layer 3 VPN route is received but the nexthop is not in the `inet.3` table, the route cannot be resolved. Because the router is not in the forwarding path, an effective workaround is to use the information in `inet.0`. The metric information in `inet.0` is useful for choosing the best route, even though it cannot be used for forwarding.

An alternative approach is to make sure that LSPs are provisioned to the RR. This happens automatically if you configure LDP.

## 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 <b>[edit]</b> hierarchy level.
Route Reflector	<pre> set routing-options resolution rib bgp.l3vpn.0 resolution-ribs inet.0 set routing-options resolution rib inet.0 import POLICY-limit-resolve-routes set policy-options policy-statement POLICY-limit-resolve-routes term isis from protocol isis set policy-options policy-statement POLICY-limit-resolve-routes term isis then accept set policy-options policy-statement POLICY-limit-resolve-routes then reject </pre>

**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 route resolution:

1. Configure `bgp.l3vpn.0` to use the information in `inet.0` instead of `inet.3` to resolve routes.

```
[edit routing-options resolution rib bgp.l3vpn.0]
user@RR# set resolution-ribs inet.0
```

2. (Optional) Configure the routing policy.

```
[edit policy-options policy-statement POLICY-limit-resolve-routes]
user@RR# set term isis from protocol isis
user@RR# set term isis then accept
user@RR# set then reject
```

3. (Optional) Apply the policy.

```
[edit routing-options resolution rib inet.0]
user@RR# set import POLICY-limit-resolve-routes
```

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

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

## Results

Confirm your configuration by issuing the `show policy-options` and `show routing-options` commands.

```
user@RR# show policy-options
policy-statement POLICY-limit-resolve-routes {
  term isis {
    from protocol isis;
    then accept;
  }
  then reject;
}

user@RR# show routing-options
resolution {
  rib bgp.l3vpn.0 {
    resolution-ribs inet.0;
  }
  rib inet.0 {
    import POLICY-limit-resolve-routes;
  }
}
```

## Verification

Confirm that the configuration is working properly by running the following commands:

- `show route`
- `show route forwarding-table`
- `show route resolution`

### Related Documentation

- [Example: Configuring BGP Route Reflectors](#)
- [Example: Configuring Route Resolution on PE Routers on page 255](#)

## Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs

You can configure a maximum limit on the number of prefixes and paths that can be installed into the routing tables. Using prefix and path limits, you can curtail the number of prefixes and paths received from a CE router in a VPN. Prefix and path limits apply only to dynamic routing protocols, and are not applicable to static or interface routes.

To limit the number of paths accepted by a PE router from a CE router, include the **maximum-paths** statement:

```
maximum-paths path-limit <log-interval interval | log-only | threshold percentage>;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Specify the **log-only** option to generate warning messages only (an advisory limit). Specify the **threshold** option to generate warnings before the limit is reached. Specify the **log-interval** option to configure the minimum time interval between log messages.

There are two modes for route limits: advisory and mandatory. An advisory limit triggers warnings. A mandatory limit rejects additional routes after the limit is reached.



**NOTE:** Application of a route limit may result in unpredictable dynamic routing protocol behavior. For example, when the limit is reached and routes are rejected, BGP may not reinstall the rejected routes after the number of routes drops back below the limit. BGP sessions may need to be cleared.

---

To limit the number of prefixes accepted by a PE router from a CE router, include the **maximum-prefixes** statement:

```
maximum-prefixes prefix-limit <log-interval interval | log-only | threshold percentage>;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.



There are two modes for route limits: advisory and mandatory. An advisory limit triggers warnings. A mandatory limit rejects additional routes after the limit is reached.



**NOTE:** Application of a route limit may result in unpredictable dynamic routing protocol behavior. For example, when the limit is reached and routes are rejected, BGP may not reinstall the rejected routes after the number of routes drops back below the limit. BGP sessions may need to be cleared.

A mandatory path or prefix limit, in addition to triggering a warning message, rejects any additional paths or prefixes once the limit is reached.



**NOTE:** Setting a path or prefix limit might result in unpredictable dynamic routing protocol behavior.

You can also configure the following options for both the **maximum-paths** and **maximum-prefixes** statements:

- **log-interval**—Specify the interval at which log messages are sent. This option generates warning messages only (an advisory limit).

Specify the **log-interval** option to configure the minimum time interval between log messages.

- **log-only**—Generate warning messages only. No limit is placed on the number of paths or prefixes stored in the routing tables.
- **threshold**—Generate warning messages after the specified percentage of the maximum paths or prefixes has been reached.



## CHAPTER 18

# Configuring Protocol Family Qualifiers for Layer 3 VPNs

- [Example: Configuring Layer 3 VPN Protocol Family Qualifiers for Route Filters on page 263](#)

### Example: Configuring Layer 3 VPN Protocol Family Qualifiers for Route Filters

---

This example shows how to control the scope of BGP import policies by configuring a family qualifier for the BGP import policy. The family qualifier specifies routes of type **inet**, **inet6**, **inet-vpn**, or **inet6-vpn**.

- [Requirements on page 263](#)
- [Overview on page 263](#)
- [Configuration on page 264](#)
- [Verification on page 266](#)

### Requirements

This example uses Junos OS Release 10.0 or later.

Before you begin:

- Configure the device interfaces.
- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library*.
- Configure a BGP session for multiple route types. For example, configure the session for both family **inet** routes and family **inet-vpn** routes. See “[Configuring IBGP Sessions Between PE Routers in VPNs](#)” on page 19 and “[Configuring Layer 3 VPNs to Carry IPv6 Traffic](#)” on page 227.

### Overview

Family qualifiers cause a route filter to match only one specific family. When you configure an IPv4 route filter without a family qualifier, as shown here, the route filter matches **inet** and **inet-vpn** routes.

```
route-filter ipv4-address/mask;
```

Likewise, when you configure an IPv6 route filter without a family qualifier, as shown here, the route filter matches **inet6** and **inet6-vpn** routes.

```
route-filter ipv6-address/mask;
```

Consider the case in which a BGP session has been configured for both family **inet** routes and family **inet-vpn** routes, and an import policy has been configured for this BGP session. This means that both family **inet** and family **inet-vpn** routes, when received, share the same import policy. The policy term might look as follows:

```
from {  
    route-filter 0.0.0.0/0 exact;  
}  
then {  
    next-hop self;  
    accept;  
}
```

This route-filter logic matches an **inet** route of 0.0.0.0 and an **inet-vpn** route whose IPv4 address portion is 0.0.0.0. The 8-byte route distinguisher portion of the **inet-vpn** route is not considered in the route-filter matching. This is a change in Junos OS behavior that was introduced in Junos OS Release 10.0.

If you do not want your policy to match both types of routes, add a family qualifier to your policy. To have the route-filter match only **inet** routes, add the family **inet** policy qualifier. To have the route-filter match only **inet-vpn** routes, add the family **inet-vpn** policy qualifier.

The family qualifier is evaluated before the route-filter is evaluated. Thus, the route-filter is not evaluated if the family match fails. The same logic applies to family **inet6** and family **inet6-vpn**. The route-filter used in the **inet6** example must use an IPv6 address. There is a potential efficiency gain in using a family qualifier because the family qualifier is tested before most other qualifiers, quickly eliminating routes from undesired families.

## Configuration

<b>CLI Quick Configuration</b>	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 <b>[edit]</b> hierarchy level.
--------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<b>inet Example</b>	<pre>set policy-options policy-statement specific-family from family inet set policy-options policy-statement specific-family from route-filter 0.0.0.0/0 exact set policy-options policy-statement specific-family then next-hop self set policy-options policy-statement specific-family then accept set protocols bgp import specific-family</pre>
---------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<b>Inet-vpn Example</b>	<pre>set policy-options policy-statement specific-family from family inet-vpn set policy-options policy-statement specific-family from route-filter 0.0.0.0/0 exact set policy-options policy-statement specific-family then next-hop self set policy-options policy-statement specific-family then accept set protocols bgp import specific-family</pre>
-------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**inet6 Example**

```
set policy-options policy-statement specific-family from family inet6
set policy-options policy-statement specific-family from route-filter 0::0/0 exact
set policy-options policy-statement specific-family then next-hop self
set policy-options policy-statement specific-family then accept
set protocols bgp import specific-family
```

**Inet6-vpn Example**

```
set policy-options policy-statement specific-family from family inet6-vpn
set policy-options policy-statement specific-family from route-filter 0::0/0 exact
set policy-options policy-statement specific-family then next-hop self
set policy-options policy-statement specific-family then accept
set protocols bgp import specific-family
```

**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 flow map:

1. Configure the family qualifier.

```
[edit policy-options]
user@host# set policy-statement specific-family from family inet
```

2. Configure the route filter.

```
[edit policy-options]
user@host# set policy-statement specific-family from route-filter 0.0.0.0/0 exact
```

3. Configure the policy actions.

```
[edit policy-options]
user@host# set policy-statement specific-family then next-hop self
user@host# set policy-statement specific-family then accept
```

4. Apply the policy.

```
[edit protocols bgp]
user@host# set import specific-family
```

## Results

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

```
user@host# show protocols
bgp {
  import specific-family;
}
user@host# show policy-options
policy-statement specific-family {
```

```
from {  
    family inet;  
    route-filter 0.0.0.0/0 exact;  
}  
then {  
    next-hop self;  
    accept;  
}  
}
```

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

Repeat the procedure for every protocol family for which you need a specific route-filter policy.

## Verification

To verify the configuration, run the following commands:

- **show route advertising-protocol bgp *neighbor* detail**
- **show route instance *instance-name* detail**

### Related Documentation

- *Understanding Route Filters for Use in Routing Policy Match Conditions*
- *Route Filter Match Conditions*
- *Example: Configuring Policy Chains and Route Filters*
- *Example: Configuring the MED Using Route Filters*
- *Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF*

# Configuring Class of Service for Layer 3 VPNs

- [Configuring Traffic Policing in Layer 3 VPNs on page 267](#)
- [Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs on page 268](#)

## Configuring Traffic Policing in Layer 3 VPNs

---

You can use policing to control the amount of traffic flowing over the interfaces servicing a Layer 3 VPN. If policing is disabled on an interface, all the available bandwidth on a Layer 3 VPN tunnel can be used by a single CCC or TCC interface.

For more information about the **policer** statement, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

To enable Layer 3 VPN policing on an interface, include the **policer** statement:

```
policer {  
    input policer-template-name;  
    output policer-template-name;  
}
```

If you configure CCC encapsulation, you can include the **policer** statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family ccc]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family ccc]

If you configure TCC encapsulation, you can include the **policer** statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family tcc]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]

### Related Documentation

- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*

## Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs

When you include the **vrf-table-label** statement in the configuration for a routing instance (as described in “[Filtering Packets in Layer 3 VPNs Based on IP Headers](#)” on page 159) but do not explicitly apply a classifier to the routing instance, the default MPLS EXP classifier is applied.

For PICs that are installed on Enhanced FPCs, you can apply a custom classifier to override the default MPLS EXP classifier for the routing instance. For detailed instructions, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*. The following instructions serve as a summary:

1. Filter traffic based on the IP header by including the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy level:

```
[edit routing-instances routing-instance-name]
vrf-table-label;
```

2. Configure a custom MPLS EXP classifier by including the appropriate statements at the **[edit class-of-service]** hierarchy level. For instructions, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

3. Configure the routing instance for CoS by including the **routing-instances** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
routing-instances routing-instance-name {
  classifiers {
    exp (classifier-name | default);
  }
}
```

4. Configure the routing instance to use the custom MPLS EXP classifier by including the **classifiers** statement at the **[edit class-of-service routing-instances routing-instance-name]** hierarchy level:

```
[edit class-of-service routing-instances routing-instance-name]
classifiers {
  exp classifier-name;
}
```

To display the MPLS EXP classifiers associated with all routing instances, issue the **show class-of-service routing-instances** command.



**NOTE:** The following caveats apply to custom MPLS EXP classifiers for routing instances:

- An Enhanced FPC is required.
- Logical systems are not supported.



- Related Documentation**
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*



# Enabling Internet Access for Layer 3 VPNs

- [Non-VRF Internet Access Through Layer 3 VPNs on page 271](#)
- [Distributed Internet Access Through Layer 3 VPNs on page 272](#)
- [Routing VPN and Internet Traffic Through Different Interfaces for Layer 3 VPNs on page 273](#)
- [Routing VPN and Outgoing Internet Traffic Through the Same Interface and Routing Return Internet Traffic Through a Different Interface on page 279](#)
- [Routing VPN and Internet Traffic Through the Same Interface Bidirectionally \(VPN Has Public Addresses\) on page 280](#)
- [Routing VPN and Internet Traffic Through the Same Interface Bidirectionally \(VPN Has Private Addresses\) on page 284](#)
- [Routing Internet Traffic Through a Separate NAT Device on page 288](#)
- [Centralized Internet Access Through Layer 3 VPNs on page 296](#)

## Non-VRF Internet Access Through Layer 3 VPNs

---

Junos OS supports Internet access from a Layer 3 virtual private network (VPN). This chapter provides examples that demonstrate how to configure a provider edge (PE) router to provide Internet access to customer edge (CE) routers in a VPN. The method you use depends on the needs and specifications of the individual network. To provide Internet access through a Layer 3 VPN, you need to configure policies on the PE router. You also need to configure the **next-table** statement at the **[edit routing-instances routing-instance-name routing-options static route]** hierarchy level. When configured, this statement can point a default route from the VPN table (routing instance) to the main routing table (default instance) inet.0. The main routing table stores all Internet routes and is where final route resolution occurs.


The following sections describe ways to provide Internet access to a CE router in a Layer 3 VPN without using the VPN routing and forwarding (VRF) interface. Because these methods effectively bypass the Layer 3 VPN, they are not discussed in detail.

- [CE Router Accesses Internet Independently of the PE Router on page 272](#)
- [PE Router Provides Layer 2 Internet Service on page 272](#)

## CE Router Accesses Internet Independently of the PE Router

In this configuration, the PE router does not provide the Internet access. The CE router sends Internet traffic either to another service provider, or to the same service provider but a different router. The PE router handles Layer 3 VPN traffic only (see [Figure 30 on page 272](#)).


*Figure 30: PE Router Does Not Provide Internet Access*



## PE Router Provides Layer 2 Internet Service

In this configuration, the PE router acts as a Layer 2 device, providing a Layer 2 connection (such as circuit cross-connect [CCC]) to another router that has a full set of Internet routes. The CE router can use just one physical interface and two logical interfaces to the PE router, or it can use multiple physical interfaces to the PE router (see [Figure 31 on page 272](#)).

*Figure 31: PE Router Connects to a Router Connected to the Internet*



## Distributed Internet Access Through Layer 3 VPNs

In this scenario, the PE routers provide Internet access to the CE routers. In the examples that follow, it is assumed that the Internet routes (or defaults) are present in the inet.0 table of the PE routers that provide Internet access to selected CE routers.


When accessing the Internet from a VPN, Network Address Translation (NAT) must be performed between the VPN's private addresses and the public addresses used on the Internet unless the VPN is using the public address space. This section includes several examples of how to provide Internet access for VPNs, most of which require that the CE routers perform the address translation. The ["Routing Internet Traffic Through a Separate NAT Device" on page 288](#) example, however, requires that the service provider supply the NAT functionality using a NAT device connected to the PE router.

In all of the examples, the VPN's public IP address pool (whose entries correspond to the translated private addresses) must be added to the inet.0 table and propagated to the Internet routers to receive reverse traffic from public destinations.

## Routing VPN and Internet Traffic Through Different Interfaces for Layer 3 VPNs

In this example, VPN and Internet traffic are routed through different interfaces. The CE router sends the VPN traffic through the VPN interface and sends the Internet traffic through a separate interface that is part of the main routing table on Router PE1 (the CE router can use either one physical interface with two logical units or two physical interfaces). NAT also occurs on the CE router (see [Figure 32 on page 273](#)).

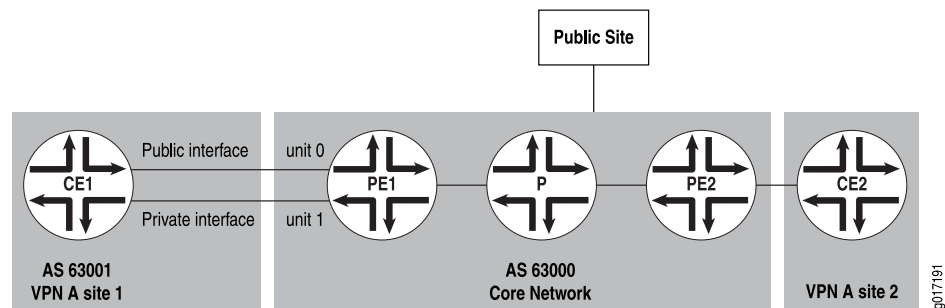
**Figure 32: Routing VPN and Internet Traffic Through Different Interfaces**



The PE router is configured to install and advertise the public IP address pool for the VPN to other core routers (for return traffic). The VPN traffic is routed normally.

[Figure 33 on page 273](#) illustrates the PE router's VPN configuration.

**Figure 33: Example of Internet Traffic Routed Through Separate Interfaces**



The configuration in this example has the following features:

- Router PE1 uses two logical interfaces to connect to Router CE1 using Frame Relay encapsulation.
- The routing protocol between Router PE1 and Router CE1 is EBGp.
- Router CE1's public IP address pool is 10.12.1.1 through 10.12.1.254 (10.12.1.0/24).
- The **next-hop-self** setting is derived from the **fix-nh policy** statement on Router PE1. PE routers are forced to use **next-hop-self** so that next-hop resolution is done only for the PE router's loopback address for non-VPN routes (by default, VPN–Internet Protocol version 4 [IPv4] routes are sent by means of **next-hop-self**).

You can configure Router CE1 with a static default route pointing to its public interface for everything else.

The following sections show how to route VPN and Internet traffic through different interfaces:

- [Configuring Interfaces on Router PE1 on page 274](#)
- [Configuring Routing Options on Router PE1 on page 274](#)
- [Configuring BGP, IS-IS, and LDP Protocols on Router PE1 on page 274](#)
- [Configuring a Routing Instance on Router PE1 on page 275](#)
- [Configuring Policy Options on Router PE1 on page 276](#)
- [Traffic Routed by Different Interfaces: Configuration Summarized by Router on page 277](#)

## Configuring Interfaces on Router PE1

Configure an interface to handle VPN traffic and an interface to handle Internet traffic:

```
[edit]
interfaces {
  t3-0/2/0 {
    dce;
    encapsulation frame-relay;
    unit 0 {
      description "to CE1 VPN interface";
      dlci 10;
      family inet {
        address 192.168.197.13/30;
      }
    }
    unit 1 {
      description "to CE1 public interface";
      dlci 20;
      family inet {
        address 192.168.198.201/30;
      }
    }
  }
}
```

## Configuring Routing Options on Router PE1

Configure a static route on Router PE1 to install a route to the CE router's public IP address pool in inet.0:

```
[edit]
routing-options {
  static {
    route 10.12.1.0/24 next-hop 192.168.198.202;
  }
}
```

## Configuring BGP, IS-IS, and LDP Protocols on Router PE1

Configure BGP on Router PE1 to allow non-VPN and VPN peering and to advertise the VPN's public IP address pool:

```
[edit]
```

```

protocols {
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.171;
      family inet {
        any;
      }
      family inet-vpn {
        any;
      }
      export [fix-nh redist-static];
      neighbor 10.255.14.177;
      neighbor 10.255.14.179;
    }
  }
}

```

Configure IS-IS on Router PE1 to allow access to internal routes:

```

[edit protocols]
isis {
  level 1 disable;
  interface so-0/0/0.0;
  interface lo0.0;
}

```

Configure LDP on Router PE1 to tunnel VPN routes:

```

[edit protocols]
ldp {
  interface so-0/0/0.0;
}

```

## Configuring a Routing Instance on Router PE1

Configure a routing instance on Router PE1:

```

[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group to-CE1 {
          peer-as 63001;
          neighbor 192.168.197.14;
        }
      }
    }
  }
}

```

## Configuring Policy Options on Router PE1

You need to configure policy options on Router PE1. The **fix-nh** policy statement sets **next-hop-self** for all non-VPN routes:

```
[edit]
policy-options {
  policy-statement fix-nh {
    then {
      next-hop self;
    }
  }
}
```

The **redist-static** policy statement advertises the VPN's public IP address pool:

```
[edit policy-options]
policy-statement redist-static {
  term a {
    from {
      protocol static;
      route-filter 10.12.1.0/24 exact;
    }
    then accept;
  }
  term b {
    then reject;
  }
}
```

Configure import and export policies for **vpna**:

```
[edit policy-options]
policy-statement vpna-import {
  term a {
    from {
      protocol bgp;
      community vpna-comm;
    }
    then accept;
  }
  term b {
    then reject;
  }
}
policy-statement vpna-export {
  term a {
    from protocol bgp;
    then {
      community add vpna-comm;
      accept;
    }
  }
  term b {
    then reject;
  }
}
```



```

}
community vpnna-comm members target:63000:100;

```

## Traffic Routed by Different Interfaces: Configuration Summarized by Router

### Router PE1

<b>Interfaces</b>	<pre> interfaces {   t3-0/2/0 {     dce;     encapsulation frame-relay;     unit 0 {       description "to CE1 VPN interface";       dlci 10;       family inet {         address 192.168.197.13/30;       }     }     unit 1 {       description "to CE1 public interface";       dlci 20;       family inet {         address 192.168.198.201/30;       }     }   } } </pre>
<b>Routing Options</b>	<pre> routing-options {   static {     route 10.12.1.0/24 next-hop 192.168.198.202;   } } </pre>
<b>BGP Protocol</b>	<pre> protocols {   bgp {     group pe-pe {       type internal;       local-address 10.255.14.171;       family inet {         any;       }       family inet-vpn {         any;       }       export [ fix-nh redist-static];       neighbor 10.255.14.177;       neighbor 10.255.14.179;     }   } } </pre>
<b>IS-IS Protocol</b>	<pre> isis {   level 1 disable; } </pre>

	<pre> interface so-0/0/0.0; interface lo0.0; } </pre>
<b>LDP Protocol</b>	<pre> ldp {   interface so-0/0/0.0; } </pre>
<b>Routing Instance</b>	<pre> routing-instances {   vpna {     instance-type vrf;     interface t3-0/2/0.0;     route-distinguisher 10.255.14.171:100;     vrf-import vpna-import;     vrf-export vpna-export;     protocols {       bgp {         group to-CE1 {           peer-as 63001;           neighbor 192.168.197.14;         }       }     }   } } </pre>
<b>Policy Options/Policy Statements</b>	<pre> policy-options {   policy-statement fix-nh {     then {       next-hop self;     }   }   policy-statement redist-static {     term a {       from {         protocol static;         route-filter 10.12.1.0/24 exact;       }       then accept;     }     term b {       then reject;     }   } } </pre>
<b>Import and Export Policies</b>	<pre> policy-statement vpna-import {   term a {     from {       protocol bgp;       community vpna-comm;     }     then accept;   } } </pre>

```


}
term b {
  then reject;
}
}
policy-statement vpna-export {
  term a {
    from protocol bgp;
    then {
      community add vpna-comm;
      accept;
    }
  }
  term b {
    then reject;
  }
}
community vpna-comm members target:63000:100;

```

## Routing VPN and Outgoing Internet Traffic Through the Same Interface and Routing Return Internet Traffic Through a Different Interface

In this example, the CE router sends VPN and Internet traffic through the same interface but receives return Internet traffic through a different interface. The PE router has a default route in the VRF table pointing to the main routing table inet.0. It routes the VPN public IP address pool (return Internet traffic) through a different interface in inet.0 (see [Figure 34 on page 279](#)). The CE router still performs NAT functions.

*Figure 34: VPN and Outgoing Internet Traffic Routed Through the Same Interface and Return Internet Traffic Routed Through a Different Interface*



The following section shows how to route VPN and outgoing Internet traffic through the same interface and routing return Internet traffic through a different interface:

- [Configuration for Router PE1 on page 279](#)

### Configuration for Router PE1

This example has the same configuration as Router PE1 in “[Routing VPN and Internet Traffic Through Different Interfaces for Layer 3 VPNs](#)” on page 273. It uses the topology shown in [Figure 33 on page 273](#). The default route to the VPN routing table is configured differently. At the `[edit routing-instances routing-instance-name routing-options]` hierarchy

level, you configure a default static route that is installed in `vpna.inet.0` and points to `inet.0` for resolution:


```
[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-table inet.0;
      }
    }
  }
}
protocols {
  bgp {
    group to-CE1 {
      peer-as 63001;
      neighbor 192.168.197.14;
    }
  }
}
```

You also need to change the configuration of Router CE1 (from the configuration that works with the configuration for Router PE1 described in [“Routing VPN and Internet Traffic Through Different Interfaces for Layer 3 VPNs” on page 273](#)) to account for the differences in the configuration of the PE routers.

## Routing VPN and Internet Traffic Through the Same Interface Bidirectionally (VPN Has Public Addresses)

This section shows how to configure a single logical interface to handle VPN and Internet traffic traveling both to and from the Internet and the CE router. This interface can handle both VPN and Internet traffic as long as there are no private addresses in the VPN. The VPN routes received from the CE router are added to the main routing table `inet.0` by means of routing table groups. This allows the PE router to attract the return traffic from the Internet (see [Figure 35 on page 280](#)).

**Figure 35: Interface Configured to Carry Both Internet and VPN Traffic**



In this example, the CE router does not need to perform NAT, because all the VPN routes are public. The CE router has a single interface to the PE router, to which it advertises VPN routes. The PE router has a default route in the VRF table pointing to the main routing

table inet.0. The PE router also imports VPN routes received from the CE router into inet.0 by means of routing table groups.

The following configuration for Router PE1 uses the same topology as in “[Routing VPN and Internet Traffic Through Different Interfaces for Layer 3 VPNs](#)” on page 273. This configuration uses a single logical interface (instead of two) between Router PE1 and Router CE1.

The following sections show how to route VPN and Internet traffic through the same interface bidirectionally (VPN has public addresses):

- [Configuring Routing Options on Router PE1 on page 281](#)
- [Configuring Routing Protocols on Router PE1 on page 281](#)
- [Configuring the Routing Instance on Router PE1 on page 282](#)
- [Traffic Routed Through the Same Interface Bidirectionally: Configuration Summarized by Router on page 283](#)

## Configuring Routing Options on Router PE1

Configure a routing table group definition for installing VPN routes in routing table groups vpna.inet.0 and inet.0:

```
[edit]
routing-options {
  rib-groups {
    vpna-to-inet0 {
      import-rib [ vpna.inet.0 inet.0 ];
    }
  }
}
```

## Configuring Routing Protocols on Router PE1

Configure MPLS, BGP, IS-IS, and LDP protocols on Router PE1. This configuration does not include the **policy redist-static** statement at the **[edit protocols bgp group pe-pe]** hierarchy level. The VPN routes are sent directly to IBGP.

Configure BGP on Router PE1 to allow non-VPN and VPN peering, and to advertise the VPN's public IP address pool:

```
[edit]
protocols {
  mpls {
    interface so-0/0/0.0;
  }
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.171;
      family inet {
        any;
      }
      family inet-vpn {
```

```

        any;
    }
    export fix-nh;
    neighbor 10.255.14.177;
    neighbor 10.255.14.173;
}
}
isis {
    level 1 disable;
    interface so-0/0/0.0;
    interface lo0.0;
}
ldp {
    interface so-0/0/0.0;
}
}

```

## Configuring the Routing Instance on Router PE1

This section describes how to configure the routing instance on Router PE1. The static route defined in the **routing-options** statement directs Internet traffic from the CE router to the inet.0 routing table. The routing table group defined by the **rib-group vpna-to-inet0** statement adds the VPN routes to inet.0.

Configure the routing instance on Router PE1:

```

[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-table inet.0;
      }
    }
    protocols {
      bgp {
        group to-CE1 {
          family inet {
            unicast {
              rib-group vpna-to-inet0;
            }
          }
        }
        peer-as 63001;
        neighbor 192.168.197.14;
      }
    }
  }
}
}

```

You must configure Router CE1 to forward all traffic to Router PE1 using a default route. Alternatively, the default route can be advertised from Router PE1 to Router CE1 with EBGp.

## Traffic Routed Through the Same Interface Bidirectionally: Configuration Summarized by Router

### Router PE1

This example uses the same configuration as in [“Routing VPN and Internet Traffic Through Different Interfaces for Layer 3 VPNs” on page 273](#). This configuration uses a single logical interface (instead of two) between Router PE1 and Router CE1.

Routing Options	<pre> routing-options {   rib-groups {     vpna-to-inet0 {       import-rib [ vpna.inet.0 inet.0 ];     }   } } </pre>
Routing Protocols	<pre> protocols {   mpls {     interface so-0/0/0.0;   }   bgp {     group pe-pe {       type internal;       local-address 10.255.14.171;       family inet {         any;       }       family inet-vpn {         any;       }       export fix-nh;       neighbor 10.255.14.177;       neighbor 10.255.14.173;     }   }   isis {     level 1 disable;     interface so-0/0/0.0;     interface lo0.0;   }   ldp {     interface so-0/0/0.0;   } } </pre>
Routing Instance	<pre> routing-instances {   vpna {     instance-type vrf;     interface t3-0/2/0.0;     route-distinguisher 10.255.14.171:100;   } } </pre>

```

vrf-import vpna-import;
vrf-export vpna-export;
routing-options {
  static {
    route 0.0.0.0/0 next-table inet.0;
  }
}
protocols {
  bgp {
    group to-CE1 {
      family inet {
        unicast {
          rib-group vpna-to-inet0;
        }
      }
    }
    peer-as 63001;
    neighbor 192.168.197.14;
  }
}
}
}
}


```

## Routing VPN and Internet Traffic Through the Same Interface Bidirectionally (VPN Has Private Addresses)

The example in this section shows how to route VPN and Internet traffic through the same interface in both directions (from the CE router to the Internet and from the Internet to the CE router). The VPN in this example has private addresses. If you can configure EBGp on the CE router, you can configure a PE router using the configuration outlined in “Routing VPN and Internet Traffic Through the Same Interface Bidirectionally (VPN Has Public Addresses)” on page 280, even if the VPN has private addresses.

In the example described in this section, the CE router uses separate communities to advertise its VPN routes and public routes. The PE router selectively imports only the public routes into the inet.0 routing table. This configuration ensures that return traffic from the Internet uses the same interface between the PE and CE routers as that used by VPN traffic going out to public Internet addresses (see Figure 36 on page 284).

**Figure 36: VPN and Internet Traffic Routed Through the Same Interface**



In this example, the CE router has one interface and a BGP session with the PE router, and it tags VPN routes and Internet routes with different communities. The PE router has one interface, selectively imports routes for the VPN's public IP address pool into inet.0, and has a default route in the VRF routing table pointing to inet.0.



The following sections show how to route VPN and Internet traffic through the same interface bidirectionally (VPN has private addresses):

- [Configuring Routing Options for Router PE1 on page 285](#)
- [Configuring a Routing Instance for Router PE1 on page 285](#)
- [Configuring Policy Options for Router PE1 on page 286](#)
- [Traffic Routed by the Same Interface Bidirectionally \(VPN Has Private Addresses\): Configuration Summarized by Router on page 286](#)

## Configuring Routing Options for Router PE1

On Router PE1, configure a routing table group to install VPN routes in the `vpna.inet.0` and `inet.0` routing tables:

```
[edit]
routing-options {
  rib-groups {
    vpna-to-inet0 {
      import-policy import-public-addr-to-inet0;
      import-rib [ vpna.inet.0 inet.0 ];
    }
  }
}
```

## Configuring a Routing Instance for Router PE1

On Router PE1, configure a routing instance. As part of the configuration for the routing instance, configure a static route that is installed in `vpna.inet.0` and is pointed at `inet.0` for resolution.

```
[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-table inet.0;
      }
    }
  }
}
```

At the `[edit routing-instances vpna protocols bgp]` hierarchy level, configure a policy (`import-public-addr-to-inet0`) to import public routes into `inet.0` and a routing table group (`vpna-to-inet0`) to allow BGP to install routes into multiple routing tables (`vpna.inet.0` and `inet.0`):

```
[edit routing-instances vpna]
protocols {
  bgp {
```

```

group to-CE1 {
  import import-public-addr-to-inet0;
  family inet {
    unicast {
      rib-group vpna-to-inet0;
    }
  }
  peer-as 63001;
  neighbor 192.168.197.14;
}
}
}

```

## Configuring Policy Options for Router PE1

Configure the policy options for Router PE1 to accept all routes initially (**term a**) and then to install routes with a **public-comm** community into routing table inet.0 (**term b**):

```

[edit]
policy-options {
  policy-statement import-public-addr-to-inet0 {
    term a {
      from {
        protocol bgp;
        rib vpna.inet.0;
        community [ public-comm private-comm ];
      }
      then accept;
    }
    term b {
      from {
        protocol bgp;
        community public-comm;
      }
      to rib inet.0;
      then accept;
    }
    term c {
      then reject;
    }
  }
  community private-comm members target:1:333;
  community public-comm members target:1:111;
  community vpna-comm members target:63000:100;
}

```

## Traffic Routed by the Same Interface Bidirectionally (VPN Has Private Addresses): Configuration Summarized by Router

### Router PE1

Routing Options	<pre> [edit] routing-options {   rib-groups {     vpna-to-inet0 {       import-policy import-public-addr-to-inet0;     }   } } </pre>
-----------------	---------------------------------------------------------------------------------------------------------------------------------------

```

        import-rib [ vpna.inet.0 inet.0 ];
    }
}

Routing Instances [edit]
routing-instances {
    vpna {
        instance-type vrf;
        interface t3-0/2/0.0;
        route-distinguisher 10.255.14.171:100;
        vrf-import vpna-import;
        vrf-export vpna-export;
        routing-options {
            static {
                route 0.0.0.0/0 next-table inet.0;
            }
        }
    }
}

Routing Instances [edit routing-instances vpna]
Protocols BGP protocols {
    bgp {
        group to-CE1 {
            family inet {
                unicast {
                    rib-group vpna-to-inet0;
                }
            }
        }
        peer-as 63001;
        neighbor 192.168.197.14;
    }
}

Policy Options [edit]
policy-options {
    policy-statement import-public-addr-to-inet0 {
        term a {
            from {
                protocol bgp;
                rib vpna.inet.0;
                community [ public-comm private-comm ];
            }
            then accept;
        }
        term b {
            from {
                protocol bgp;
                community public-comm;
            }
            to rib inet.0;
            then accept;
        }
    }
}

```

```


    }
    term c {
        then reject;
    }
}
community private-comm members target:1:333;
community public-comm members target:1:111;
community vpna-comm members target:63000:100;
}

```

## Routing Internet Traffic Through a Separate NAT Device

In this example, the CE router does not perform NAT. It sends both VPN and Internet traffic over the same interface to the PE router. The PE router is connected to an NAT device by means of two interfaces. One interface is configured in the PE router's VRF table and points to a VPN interface on the NAT device, which can route Internet traffic for the VPN. The other interface is in a default instance; for example, part of public routing table inet.0. There can be a single physical connection between the PE router and the NAT device and multiple logical connections—one for each VRF table and another interface—as part of the global routing table (see [Figure 37 on page 288](#)).

**Figure 37: Internet Traffic Routed Through a Separate NAT Device**



- [Requirements on page 288](#)
- [Overview on page 288](#)
- [Configuration on page 289](#)

## Requirements

This example uses the following hardware and software components:


- M Series routers
- Junos OS Release 9.3 or later

## Overview

This example's topology expands upon that illustrated in [Figure 33 on page 273](#). The CE router sends both VPN and Internet traffic to Router PE1. VPN traffic is routed based on the VPN routes received by Router PE1. Traffic for everything else is sent to the NAT device using Router PE1's private interface to the NAT device, which then translates the private addresses and sends the traffic back to Router PE1 using that router's public interface (see [Figure 38 on page 289](#)).

## Topology

Figure 38: Internet Traffic Routed Through a NAT Example Topology



## Configuration

To route Internet traffic through a separate NAT device, perform these tasks:

- [Configuring Interfaces on Router PE1 on page 289](#)
- [Configuring Routing Options for Router PE1 on page 290](#)
- [Configuring Routing Protocols on Router PE1 on page 290](#)
- [Configuring a Routing Instance on Router PE1 on page 291](#)
- [Results on page 293](#)

### Configuring Interfaces on Router PE1

#### Step-by-Step Procedure

1. Configure an interface for VPN traffic from Router CE1:

```
[edit]
interfaces {
  t3-0/2/0 {
    dce;
    encapsulation frame-relay;
    unit 0 {
      description "to CE1 VPN interface";
      dlci 10;
      family inet {
        address 192.168.197.13/30;
      }
    }
  }
}
```

2. Configure an interface for VPN traffic to and from the NAT device (unit 0), and an interface for Internet traffic to and from the NAT device (unit 1):

```
[edit]
interfaces {
  at-1/3/1 {
    atm-options {
      vpi 1 maximum-vc 255;
    }
    unit 0 {
      description "to NAT VPN interface";
      vci 1.100;
      family inet {
        address 10.23.0.2/32 {
          destination 10.23.0.1;
        }
      }
    }
    unit 1 {
      description "to NAT public interface";
      vci 1.101;
      family inet {
        address 10.23.0.6/32 {
          destination 10.23.0.5;
        }
      }
    }
  }
}
```

---

### Configuring Routing Options for Router PE1

#### Step-by-Step Procedure

1. Configure a static route on Router PE1 to direct Internet traffic to the CE router through the NAT device. Router PE1 distributes this route to the Internet.

```
[edit]
routing-options {
  static {
    route 10.12.1.0/24 next-hop 10.23.0.5;
  }
}
```

---

### Configuring Routing Protocols on Router PE1

#### Step-by-Step Procedure

Configure the following routing protocols on Router PE1:

1. Configure MPLS on Router PE1. Include the NAT device's VPN interface in the VRF table.

```
[edit]
protocols {
  mpls {
    interface so-0/0/0.0;
```

```

        interface at-1/3/1.0;
    }
}

```

2. Configure BGP on Router PE1. Include a policy to advertise the public IP address pool:

```

[edit]
protocols {
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.171;
      family inet {
        any;
      }
      family inet-vpn {
        any;
      }
      export [ fix-nh redist-static ];
      neighbor 10.255.14.177;
      neighbor 10.255.14.173;
    }
  }
}

```

3. Configure IS-IS on Router PE1:

```

[edit]
protocols {
  isis {
    level 1 disable;
    interface so-0/0/0.0;
    interface lo0.0;
  }
}

```

4. Configure LDP on Router PE1:

```

[edit]
protocols {
  ldp {
    interface so-0/0/0.0;
  }
}

```

### Configuring a Routing Instance on Router PE1

#### Step-by-Step Procedure

Configure the Layer 3 VPN routing instance on Router PE1:

1. Configure a routing instance on Router PE1. As part of the routing instance configuration, under **routing-options**, configure a static default route in `vpna.inet.0`

pointing to the NAT device's VPN interface (this directs all non-VPN traffic to the NAT device):

```
[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    interface at-1/3/1.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 10.23.0.1;
      }
    }
    protocols {
      bgp {
        group to-CE1 {
          peer-as 63001;
          neighbor 192.168.197.14;
        }
      }
    }
  }
}
```

2. Configure the routing policy for the Layer 3 VPN routing instance on Router PE1:

```
policy-options {
  policy-statement fix-nh {
    then {
      next-hop self;
    }
  }
  policy-statement redistrib-static {
    term a {
      from {
        protocol static;
        route-filter 10.12.1.0/24 exact;
      }
      then accept;
    }
    term b {
      from protocol bgp;
      then accept;
    }
    term c {
      then accept;
    }
  }
  policy-statement vpna-import {
    term a {
      from {
        protocol bgp;

```



```

        community vpna-comm;
    }
    then accept;
}
term b {
    then reject;
}
}
policy-statement vpna-export {
    term a {
        from protocol bgp;
        then {
            community add vpna-comm;
            accept;
        }
    }
    term b {
        then reject;
    }
}
community vpna-comm members target:63000:100;
}

```

## Results

From configuration mode on Router PE1, confirm your configuration by entering the show interfaces, show routing-options, show protocols, show routing-instances and show policy-options commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE1# show interfaces
interfaces {
    t3-0/2/0 {
        dce;
        encapsulation frame-relay;
        unit 0 {
            description "to CE1 VPN interface";
            dlci 10;
            family inet {
                address 192.168.197.13/30;
            }
        }
    }
}
at-1/3/1 {
    atm-options {
        vpi 1 maximum-vcs 255;
    }
    unit 0 {
        description "to NAT VPN interface";
        vci 1.100;
        family inet {
            address 10.23.0.2/32 {
                destination 10.23.0.1;
            }
        }
    }
}

```

```
    }
  }
  unit 1 {
    description "to NAT public interface";
    vci 1.101;
    family inet {
      address 10.23.0.6/32 {
        destination 10.23.0.5;
      }
    }
  }
}

user@PE1# show routing-options
routing-options {
  static {
    route 10.12.1.0/24 next-hop 10.23.0.5;
  }
}

user@PE1# show protocols
protocols {
  mpls {
    interface so-0/0/0.0;
    interface at-1/3/1.0;
  }
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.171;
      family inet {
        any;
      }
      family inet-vpn {
        any;
      }
      export [ fix-nh redist-static ];
      neighbor 10.255.14.177;
      neighbor 10.255.14.173;
    }
  }
  isis {
    level 1 disable;
    interface so-0/0/0.0;
    interface lo0.0;
  }
  ldp {
    interface so-0/0/0.0;
  }
}

user@PE1# show routing-instances
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
  }
}
```

```
interface at-1/3/1.0;
route-distinguisher 10.255.14.171:100;
vrf-import vpna-import;
vrf-export vpna-export;
routing-options {
  static {
    route 0.0.0.0/0 next-hop 10.23.0.1;
  }
}
protocols {
  bgp {
    group to-CE1 {
      peer-as 63001;
      neighbor 192.168.197.14;
    }
  }
}
```

user@PE1# show policy-options

```
policy-options {
  policy-statement fix-nh {
    then {
      next-hop self;
    }
  }
  policy-statement redist-static {
    term a {
      from {
        protocol static;
        route-filter 10.12.1.0/24 exact;
      }
      then accept;
    }
    term b {
      from protocol bgp;
      then accept;
    }
    term c {
      then accept;
    }
  }
  policy-statement vpna-import {
    term a {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpna-export {
```

```
term a {  
    from protocol bgp;  
    then {  
        community add vpna-comm;  
        accept;  
    }  
}  
term b {  
    then reject;  
}  
}  
community vpna-comm members target:63000:100;  
}
```

---

## Centralized Internet Access Through Layer 3 VPNs

This section describes several ways to configure a CE router to act as a central site for Internet access. Internet traffic from other sites (CE routers) is routed to the hub CE router (which also performs NAT) using that router's VPN interface. The hub CE router then forwards the traffic to a PE router connected to the Internet through another interface identified in the inet.0 table. The hub CE router can advertise a default route to the spoke CE routers. The disadvantage of this type of configuration is that all traffic has to go through the central CE router before going to the Internet, causing network delays if this router receives too much traffic. However, in a corporate network, traffic might have to be routed to a central site because most corporate networks separate the VPN from the Internet by means of a single firewall.


This section includes the following examples:

- [Routing Internet Traffic Through a Hub CE Router on page 296](#)
- [Routing Internet Traffic Through Multiple CE Routers on page 300](#)

### Routing Internet Traffic Through a Hub CE Router


In this example, Internet traffic is routed through a hub CE router. The hub CE router has two interfaces to the hub PE router: a VPN interface and a public interface. It performs NAT on traffic forwarded from the hub PE router through the VPN interface and forwards that traffic from its public interface back to the hub PE router. The hub PE router has a static default route in its VRF table pointing to the hub CE router's VPN interface. It announces this default route to the rest of the VPN, attracting all non-VPN traffic to the hub CE route. The hub PE router also installs and distributes the VPN's public IP address space (see [Figure 39 on page 297](#)).

Figure 39: Internet Access Through a Hub CE Router Performing NAT



The configuration for this example is almost identical to that described in [“Routing Internet Traffic Through a Separate NAT Device”](#) on page 288. The difference is that Router PE1 is configured to announce a static default route to the other CE routers (see [Figure 40](#) on page 297).

Figure 40: Internet Access Provided Through a Hub CE Router



The following sections show how to configure centralized Internet access by routing Internet traffic through a hub CE router:

- [Configuring a Routing Instance on Router PE1](#) on page 297
- [Configuring Policy Options on Router PE1](#) on page 298
- [Internet Traffic Routed by a Hub CE Router: Configuration Summarized by Router](#) on page 299

### Configuring a Routing Instance on Router PE1

Configure a routing instance for Router PE1. As part of this configuration, under **routing-options**, configure a default static route (**route 0.0.0.0/0**) to be installed in

vpna.inet.0, and point the route to the hub CE router's VPN interface (10.23.0.1). Also, configure BGP under the routing instance to export the default route to the local CE router:

```
[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    interface at-1/3/1.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 10.23.0.1;
      }
    }
    protocols {
      bgp {
        group to-CE1 {
          export export-default;
          peer-as 63001;
          neighbor 192.168.197.14;
        }
      }
    }
  }
}
```

### Configuring Policy Options on Router PE1

Configure policy options on Router PE1. As part of this configuration, Router PE1 should export the static default route to all the remote PE routers in **vpna** (configured in the **policy-statement vpna-export** statement under **term b**):

```
[edit]
policy-options {
  policy-statement vpna-export {
    term a {
      from protocol bgp;
      then {
        community add vpna-comm;
        accept;
      }
    }
    term b {
      from {
        protocol static;
        route-filter 0.0.0.0/0 exact;
      }
      then {
        community add vpna-comm;
        accept;
      }
    }
  }
}
```

```

        term c {
            then reject;
        }
    }
    policy-statement export-default {
        term a {
            from {
                protocol static;
                route-filter 0.0.0.0/0 exact;
            }
            then accept;
        }
        term b {
            from protocol bgp;
            then accept;
        }
        term c {
            then reject;
        }
    }
}

```

### Internet Traffic Routed by a Hub CE Router: Configuration Summarized by Router

#### Router PE1

The configuration for Router PE1 is almost identical to that for the example in [“Routing Internet Traffic Through a Separate NAT Device” on page 288](#). The difference is that Router PE1 is configured to announce a static default route to the other CE routers.

<b>Routing Instance</b>	<pre> routing-instances {     vpna {         instance-type vrf;         interface t3-0/2/0.0;         interface at-1/3/1.0;         route-distinguisher 10.255.14.171:100;         vrf-import vpna-import;         vrf-export vpna-export;         routing-options {             static {                 route 0.0.0.0/0 next-hop 10.23.0.1;             }         }         protocols {             bgp {                 group to-CE1 {                     export export-default;                     peer-as 63001;                     neighbor 192.168.197.14;                 }             }         }     } } </pre>
-------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------


<b>Policy Options</b>	<pre> policy-options {   policy-statement vpna-export {     term a {       from protocol bgp;       then {         community add vpna-comm;         accept;       }     }     term b {       from {         protocol static;         route-filter 0.0.0.0/0 exact;       }       then {         community add vpna-comm;         accept;       }     }     term c {       then reject;     }   }   policy-statement export-default {     term a {       from {         protocol static;         route-filter 0.0.0.0/0 exact;       }       then accept;     }     term b {       from protocol bgp;       then accept;     }     term c {       then reject;     }   } } </pre>
-----------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## Routing Internet Traffic Through Multiple CE Routers

The example in this section is an extension of that described in [“Centralized Internet Access Through Layer 3 VPNs” on page 296](#). This example provides different exit points for different sites by means of multiple hub CE routers that perform similar functions. Each hub CE router tags the default route with a different route target and allows the spoke CE routers to select the hub site that should be used for Internet access (see [Figure 41 on page 301](#)).



Figure 41: Two Hub CE Routers Handling Internet Traffic and NAT



This example uses two hub CE routers that handle NAT and Internet traffic:

- Hub1 CE router tags 0/0 with community **public-comm1** (target: 1:111)
- Hub2 CE router tags 0/0 with community **public-comm2** (target: 1:112)

The spoke CE router in this example is configured to have a bias toward Hub2 for Internet access.

The following sections describe how to configure two hub CE routers to handle Internet traffic and NAT:

- [Configuring a Routing Instance on Router PE1 on page 301](#)
- [Configuring Policy Options on Router PE1 on page 302](#)
- [Configuring a Routing Instance on Router PE3 on page 303](#)
- [Configuring Policy Options on Router PE3 on page 303](#)
- [Routing Internet Traffic Through Multiple CE Routers: Configuration Summarized by Router on page 304](#)

### Configuring a Routing Instance on Router PE1

Configure a routing instance on Router PE1:

```
[edit]
routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    interface at-1/3/1.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
```

```

routing-options {
  static {
    route 0.0.0.0/0 next-hop 10.23.0.1;
  }
}
protocols {
  bgp {
    group to-CE1 {
      export export-default;
      peer-as 63001;
      neighbor 192.168.197.14;
    }
  }
}
}

```

### Configuring Policy Options on Router PE1

The policy options for Router PE1 are the same as in [“Routing Internet Traffic Through a Hub CE Router” on page 296](#), but the configuration in this example includes an additional community, **public-comm1**, in the **export** statement:

```

[edit]
policy-options {
  policy-statement vpna-import {
    term a {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpna-export {
    term a {
      from {
        protocol static;
        route-filter 0.0.0.0/0 exact;
      }
      then {
        community add public-comm1;
        community add vpna-comm;
        accept;
      }
    }
    term b {
      from protocol bgp;
      then {
        community add vpna-comm;
        accept;
      }
    }
  }
}

```

```

    term c {
        then reject;
    }
}
community public-comm1 members target:1:111;
community public-comm2 members target:1:112;
community vpna-comm members target:63000:100;
}

```

The configuration of Router PE2 is identical to that of Router PE1 except that Router PE2 exports the default route through community **public-comm2**.

### Configuring a Routing Instance on Router PE3

Configure routing instance **vpna** on Router PE3:

```

[edit]
routing-instances {
    vpna {
        instance-type vrf;
        interface t1-0/2/0.0;
        route-distinguisher 10.255.14.173:100;
        vrf-import vpna-import;
        vrf-export vpna-export;
        protocols {
            rip {
                group to-vpn12 {
                    export export-CE;
                    neighbor t1-0/2/0.0;
                }
            }
        }
    }
}

```

### Configuring Policy Options on Router PE3

Configure the **vrf-import** policy for Router PE3 to select the Internet exit point based on the additional communities specified in [“Configuring Policy Options on Router PE1” on page 302](#):

```

[edit]
policy-options {
    policy-statement vpna-export {
        term a {
            from protocol rip;
            then {
                community add vpna-comm;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
    policy-statement vpna-import {

```

```

term a {
  from {
    protocol bgp;
    community public-comm1;
    route-filter 0.0.0.0/0 exact;
  }
  then reject;
}
term b {
  from {
    protocol bgp;
    community vpna-comm;
  }
  then accept;
}
term c {
  then reject;
}
}
policy-statement export-CE {
  from protocol bgp;
  then accept;
}
community vpna-comm members target:69:100;
community public-comm1 members target:1:111;
community public-comm2 members target:1:112;
}

```

### Routing Internet Traffic Through Multiple CE Routers: Configuration Summarized by Router

#### Router PE1

This configuration is an extension of the example in [“Routing Internet Traffic Through a Hub CE Router” on page 296](#). It provides different exit points for various sites by using multiple hub CE routers that perform similar functions.

#### Routing Instances

```

routing-instances {
  vpna {
    instance-type vrf;
    interface t3-0/2/0.0;
    interface at-1/3/1.0;
    route-distinguisher 10.255.14.171:100;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 10.23.0.1;
      }
    }
  }
  protocols {
    bgp {
      group to-CE1 {
        export export-default;
        peer-as 63001;
        neighbor 192.168.197.14;
      }
    }
  }
}

```

```

    }
  }
}
}

```

**Policy Options**

```

policy-options {
  policy-statement vpna-import {
    term a {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpna-export {
    term a {
      from {
        protocol static;
        route-filter 0.0.0.0/0 exact;
      }
      then {
        community add public-comm1;
        community add vpna-comm;
        accept;
      }
    }
    term b {
      from protocol bgp;
      then {
        community add vpna-comm;
        accept;
      }
    }
    term c {
      then reject;
    }
  }
  community public-comm1 members target:1:111;
  community public-comm2 members target:1:112;
  community vpna-comm members target:63000:100;
}

```

**Router PE2**

The configuration of Router PE2 is identical to that of Router PE1, except that Router PE2 exports the default route through community **public-comm2** (see [“Policy Options” on page 305](#)).

**Router PE3**

<b>Routing Instances</b>	<pre> routing-instances {   vpn {     instance-type vrf;     interface t1-0/2/0.0;     route-distinguisher 10.255.14.173:100;     vrf-import vpn-import;     vrf-export vpn-export;     protocols {       rip {         group to-vpn12 {           export export-CE;           neighbor t1-0/2/0.0;         }       }     }   } } </pre>
--------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<b>Policy Options</b>	<pre> policy-options {   policy-statement vpn-export {     term a {       from protocol rip;       then {         community add vpn-comm;         accept;       }     }     term b {       then reject;     }   }   policy-statement vpn-import {     term a {       from {         protocol bgp;         community public-comm1;         route-filter 0.0.0.0/0 exact;       }       then reject;     }     term b {       from {         protocol bgp;         community vpn-comm;       }       then accept;     }     term c {       then reject;     }   }   policy-statement export-CE { </pre>
-----------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

```
    from protocol bgp;  
    then accept;  
}  
community vpna-comm members target:69:100;  
community public-comm1 members target:1:111;  
community public-comm2 members target:1:112;  
}
```





## PART 3

# Configuring Interprovider and Carrier-of-Carrier VPNs

- [Interprovider and Carrier-of-Carrier VPNs Overview on page 311](#)
- [Configuring Interprovider VPNs on page 317](#)
- [Configuring Carrier-of-Carrier VPNs on page 399](#)



# Interprovider and Carrier-of-Carrier VPNs Overview

- [Traditional VPNs, Interprovider VPNs, and Carrier-of-Carriers VPNs on page 311](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)
- [Interprovider and Carrier-of-Carrier VPNs Example Terminology on page 313](#)
- [Supported Carrier-of-Carriers and Interprovider VPN Standards on page 313](#)
- [Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics on page 314](#)

## Traditional VPNs, Interprovider VPNs, and Carrier-of-Carriers VPNs

---

As VPNs are deployed on the Internet, the customer of a VPN service provider might be another service provider rather than an end customer. The customer service provider depends on the VPN service provider to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks.

If the customer service provider's sites have different autonomous system (AS) numbers, then the VPN transit service provider supports carrier-of-carrier VPN service for the interprovider VPN service. If the customer service provider's sites have the same AS number, then the VPN transit service provider delivers a carrier-of-carriers VPN service.

There are several different methods for enabling interprovider VPNs based on RFC 4364, **BGP/MPLS IP Virtual Private Networks (VPNs)**:

- Interprovider Layer 3 VPN Option A—Interprovider VRF-to-VRF connections at the AS boundary routers (ASBR) (not very scalable).
- Interprovider Layer 3 VPN Option B—Interprovider EBGp redistribution of labeled VPN-IPv4 routes from AS to neighboring AS (somewhat scalable).
- Interprovider Layer 3 VPN Option C—Interprovider multihop EBGp redistribution of labeled VPN-IPv4 routes between source and destination ASs, with EBGp redistribution of labeled IPv4 routes from AS to neighboring AS (very scalable).

In traditional IP routing architectures, there is a clear distinction between internal routes and external routes. From the perspective of an Internet service provider (ISP), internal routes include all the provider's internal links (including BGP next hops) and loopback

interfaces. These internal routes are exchanged with other routing platforms in the ISP's network by means of an interior gateway protocol (IGP), such as OSPF or IS-IS. All routes learned at Internet peering points or from customer sites are classified as external routes and are distributed by means of an exterior gateway protocol (EGP) such as BGP. In traditional IP routing architectures, the number of internal routes is typically much smaller than the number of external routes.

**Related  
Documentation**

- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)
- [Understanding Carrier-of-Carriers VPNs on page 400](#)
- [Interprovider VPNs on page 317](#)
- *MPLS Feature Support on QFX Series and EX4600 Switches*

---

## Understanding Interprovider and Carrier-of-Carriers VPNs

---

All interprovider and carrier-of-carriers VPNs share the following characteristics:

- Each interprovider or carrier-of-carriers VPN customer must distinguish between internal and external customer routes.
- Internal customer routes must be maintained by the VPN service provider in its PE routers.
- External customer routes are carried only by the customer's routing platforms, not by the VPN service provider's routing platforms.

The key difference between interprovider and carrier-of-carriers VPNs is whether the customer sites belong to the same AS or to separate ASs:

- ["Interprovider VPNs" on page 317](#)—The customer sites belong to different ASs. You need to configure EBGp to exchange the customer's external routes.
- ["Understanding Carrier-of-Carriers VPNs" on page 400](#)—The customer sites belong to the same AS. You need to configure IBGP to exchange the customer's external routes.

In general, each service provider in a VPN hierarchy is required to maintain its own internal routes in its P routers, and the internal routes of its customers in its PE routers. By recursively applying this rule, it is possible to create a hierarchy of VPNs.

The following are definitions of the types of PE routers specific to interprovider and carrier-of-carriers VPNs:

- The AS border router is located at the AS border and handles traffic leaving and entering the AS.
- The end PE router is the PE router in the customer VPN; it is connected to the CE router at the end customer's site.

**Related  
Documentation**

- [Understanding Carrier-of-Carriers VPNs on page 400](#)
- [Interprovider VPNs on page 317](#)

- *MPLS Feature Support on QFX Series and EX4600 Switches*

## Interprovider and Carrier-of-Carrier VPNs Example Terminology

### B

**bgp.l3vpn.0** The table on the provider edge (PE) router in which the VPN-IPv4 routes that are received from another PE router are stored. Incoming routes are checked against the **vrf-import** statements from all the VPNs configured on the PE router. If there is a match, the VPN-Internet Protocol version 4 (IPv4) route is added to the bgp.l3vpn.0 table. To view the bgp.l3vpn.0 table, issue the **show route table bgp.l3vpn.0** command.

### M

**MP-EBGP** The multiprotocol external BGP (MP-EBGP) mechanism is used to export VPN-IPv4 routes across an autonomous system (AS) boundary. To apply this mechanism, use the **labeled-unicast** statement at the **[edit protocols bgp group group-name family inet]** hierarchy level.

### R

**routing-instance-name.  
inet.0**

- The routing table for a specific routing instance. For example, a routing instance called VPN-A has a routing table called VPN-A.inet.0. Routes are added to this table in the following ways:
  - They are sent from a customer edge (CE) router configured within the VPN-A routing instance.
  - They are advertised from a remote PE router that passes the **vrf-import** policy configured within VPN-A (to view the route, run the **show route** command). IPv4 (not VPN-IPv4) routes are stored in this table.

### V

**vrf-export policy-name** An export policy configured on a particular routing instance on a PE router. It is required for the configuration of interprovider and carrier-of-carriers VPNs. It is applied to VPN-IPv4 routes (originally learned from locally connected CE routers as IPv4 routes), which are advertised to another PE router or route reflector.

**vrf-import policy-name** An import policy configured on a particular routing instance on a PE router. This policy is required for the configuration of interprovider and carrier-of-carriers VPNs. It is applied to VPN-IPv4 routes learned from another PE router or a route reflector.

## Supported Carrier-of-Carriers and Interprovider VPN Standards

Junos OS substantially supports the following RFCs, which define standards for carrier-of-carriers and interprovider virtual private networks (VPNs).

- RFC 3107, *Carrying Label Information in BGP-4*
- RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

Supported on MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP.

- RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture*

Supported on MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP.

- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 5601, *Pseudowire (PW) Management Information Base (MIB)*
- RFC 5603, *Ethernet Pseudowire (PW) Management Information Base (MIB)*
- RFC 6368, *Internal BGP as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

#### Related Documentation

- *Supported VPWS Standards*
- *Supported Layer 2 VPN Standards*
- [Supported Layer 3 VPN Standards on page 130](#)
- [Supported Multicast VPN Standards on page 473](#)
- *Supported VPLS Standards*
- *Supported Standards for BGP*
- *Accessing Standards Documents on the Internet*

---

## Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics

You can configure BGP to gather traffic statistics for interprovider and carrier-of-carriers VPNs.

To configure BGP to gather traffic statistics for interprovider and carrier-of-carriers VPNs, include the **traffic-statistics** statement:

```
traffic-statistics {  
    file filename <world-readable | no-world-readable>;  
    interval seconds;  
}
```

For a list of the hierarchy levels at which you can include this statement, see the summary section for this statement.



**NOTE:** Traffic statistics for interprovider and carrier-of-carriers VPNs are available only for IPv4. IPv6 is not supported.

If you do not specify a filename, the statistics are not written to a file. However, if you have included the **traffic-statistics** statement in the BGP configuration, the statistics are still available and can be accessed by means of the **show bgp group traffic-statistics group-name** command.

To account for traffic from each customer separately, separate labels must be advertised for the same prefix to the peer routers in different groups. To enable separate traffic accounting, you need to include the **per-group-label** statement in the configuration for each BGP group. By including this statement, statistics are collected and displayed that account for traffic sent by the peers of the specified BGP group.

If you configure the statement at the **[edit protocols bgp family inet]** hierarchy level, rather than configuring it for a specific BGP group, then the traffic statistics are shared with all BGP groups configured with the **traffic-statistics** statement but not configured with the **per-group-label** statement.

To account for traffic from each customer separately, include the **per-group-label** statement in the configuration for each BGP group:

```
per-group-label;
```

For a list of the hierarchy levels at which you can include this statement, see the summary section for this statement.

The following shows a sample of the output to the traffic statistics file:

```
Dec 19 10:39:54 Statistics for BGP group ext2 (Index 1) NLRI inet-labeled-unicast
Dec 19 10:39:54  FEC                Packets        Bytes    EgressAS  FECLabel
Dec 19 10:39:54  10.255.245.55          0           0         I       100160
Dec 19 10:39:54  10.255.245.57          0           0         I       100112
Dec 19 10:39:54  192.0.2.1              0           0        25       100080
Dec 19 10:39:54  192.0.2.2              0           0        25       100080
Dec 19 10:39:54  192.0.2.3             109         9592        25       100048
Dec 19 10:39:54  192.0.2.4             109         9592        25       100048
Dec 19 10:39:54  192.168.25.0           0           0         I       100064
Dec 19 10:39:54  Dec 19 10:39:54, read statistics for 5 FECs in 00:00:00 seconds
(10 queries) for BGP group ext2 (Index 1) NLRI inet-labeled-unicast
```

#### Related Documentation

- *MPLS Feature Support on QFX Series and EX4600 Switches*
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)





## CHAPTER 22


# Configuring Interprovider VPNs

- [Interprovider VPNs on page 317](#)
- [Interprovider VPN Example—MP-EBGP Between ISP Peer Routers on page 319](#)
- [Interprovider VPN Example—Multihop MP-EBGP with P Routers on page 326](#)
- [Example: Configuring Interprovider Layer 3 VPN Option A on page 334](#)
- [Example: Configuring Interprovider Layer 3 VPN Option B on page 355](#)
- [Example: Configuring Interprovider Layer 3 VPN Option C on page 375](#)

## Interprovider VPNs

Interprovider VPNs provide connectivity between separate ASs. This functionality might be used by a VPN customer who has connections to several different service providers, or different connections to the same service provider in different geographic regions, each of which has a different AS. [Figure 42 on page 317](#) illustrates the type of network topology used by an interprovider VPN.

*Figure 42: Interprovider VPN Network Topology*



The following sections describe the ways you can configure an interprovider VPN:

- [Linking VRF Tables Between Autonomous Systems on page 318](#)
- [Configuring Next Generation Layer 3 VPNs Options A, B, and C on page 318](#)
- [Configuring Multihop MP-EBGP Between AS Border Routers on page 319](#)

## Linking VRF Tables Between Autonomous Systems

You can connect two separate ASs by simply linking the VPN routing and forwarding (VRF) table in the AS border router (ASBR) of one AS to the VRF table in the ASBR in the other AS. Each ASBR must include a VRF routing instance for each VPN configured in both service provider networks. You then configure an IP session between the two ASBRs. In effect, the ASBRs treat each other as customer edge (CE) routers.

Because of the complexity of the configuration, particularly with regard to scaling, this method is not recommended. The details of this configuration are not provided with documentation.

## Configuring Next Generation Layer 3 VPNs Options A, B, and C

For next generation Layer 3 VPNs, the PE routers within an AS use multiprotocol external BGP (MP-EBGP) to distribute labeled VPN–Internet Protocol version 4 (IPv4) routes to an ASBR or to a route reflector of which the ASBR is a client. The ASBR uses multiprotocol external BGP (MP-EBGP) to distribute the labeled VPN-IPv4 routes to its peer ASBR in the neighboring AS. The peer ASBR then uses MP-IBGP to distribute labeled VPN-IPv4 routes to PE routers, or to a route reflector of which the PE routers are a client.

You can configure both unicast (Junos OS Release 9.5 and later) and multicast (Junos OS Release 12.1 and later) next generation Layer 3 VPNs across ASs. The Junos OS software supports next generation Layer 3 VPNs option A, option B, and option C:

- Option A—This is simple though less scalable interprovider VPN solution to the problem of providing VPN services to a customer that has different sites, not all of which can use the same service provider. In this implementation, the VPN routing and forwarding (VRF) table in the ASBR of one AS is linked to the VRF table in the ASBR in the other AS. Each ASBR must include a VRF instance for each VPN configured in both service provider networks. Then an IGP or BGP must be configured between the ASBRs.
- Option B—For this interprovider VPN solution, the customer requires VPN services for different sites, yet the same service provider is not available for all of those sites. With option B, the ASBR routers keep all VPN-IPv4 routes in the routing information base (RIB), and the labels associated with the prefixes are kept in the forwarding information base (FIB). Because the RIB and FIB tables can take too much of the respective allocated memory, this solution is not very scalable for an interprovider VPN. If a transit service provider is used between service provider 1 and service provider 2, the transit service provider also has to keep all VPN-IPv4 routes in the RIB and the corresponding labels in the FIB. The ASBRs at the transit service provider have the same functionality as ASBRs at service provider 1 or service provider 2 in this solution. The PE routers within each AS use multiprotocol internal BGP (MP-IBGP) to distribute labeled VPN-IPv4 routes to an ASBR or to a route reflector of which the ASBR is a client. The ASBR uses MP-EBGP to distribute the labeled VPN-IPv4 routes to its peer ASBR router in the neighboring AS. The peer ASBR then uses MP-IBGP to distribute labeled VPN-IPv4 routes to PE routers, or to a route reflector of which the PE routers are a client.
- Option C—For this interprovider VPN solution, the customer service provider depends on the VPN service provider to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks. This functionality

might be used by a VPN customer who has connections to several different service providers, or different connections to the same service provider in different geographic regions, each of which has a different AS number. For option C, only routes internal to the service provider networks are announced between ASBRs. This is achieved by using the **family inet labeled-unicast** statements in the IBGP and EBGP configuration on the PE routers. Labeled IPv4 (not VPN-IPv4) routes are exchanged by the ASBRs to support MPLS. An MP-EBGP session between the end PE routers is used for the announcement of VPN-IPv4 routes. In this manner, VPN connectivity is provided while keeping VPN-IPv4 routes out of the core network.

#### Related Documentation

- [Example: Configuring Interprovider Layer 3 VPN Option A on page 334](#)
- [Example: Configuring Interprovider Layer 3 VPN Option B on page 355](#)
- [Example: Configuring Interprovider Layer 3 VPN Option C on page 375](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)
- [Understanding Carrier-of-Carriers VPNs on page 400](#)

## Configuring Multihop MP-EBGP Between AS Border Routers

In this type of interprovider VPN configuration, P routers do not need to store all the routes in all the VPNs. Only the PE routers must have all the VPN routes. The P routers simply forward traffic to the PE routers—they do not store or process any information about the packets' destination. The connections between the AS border routers in separate ASs forward traffic between the ASs, much as a label-switched path (LSP) works.

The following are the basic steps you take to configure an interprovider VPN in this manner:

1. Configure multihop EBGP redistribution of labeled VPN-IPv4 routes between the source and destination ASs.
2. Configure EBGP to redistribute labeled IPv4 routes from its AS to neighboring ASs.
3. Configure MPLS on the end PE routers of the VPNs.

## Interprovider VPN Example—MP-EBGP Between ISP Peer Routers

In this example, all routes learned from the CE routers (or switches) are sent over both service provider networks as VPN-IPv4 routes. The routes are initially learned by the PE routers (Router B and Router E) from the CE routers (Router A and Router F) and are announced by the PE routers to the AS border routers (Router C and Router D). The AS border routers are then configured with an MP-EBGP session, enabling them to pass the VPN-IPv4 routes with each other. When an AS border router—Router C for example—learns VPN-IPv4 routes from an IBGP PE, the following events occur:


1. Router C sets itself as the next hop for the route and creates a label for that route.
2. Router C advertises the VPN-IPv4 route to PE Router D in AS 10045.

- Router D sets the next hop to itself, creates another label, and then forwards the label and the route to its IBGP PE router (Router E).

This example has scaling limitations because of restrictions on the number of labels each PE router needs to allocate at the AS border.

Figure 43 on page 320 illustrates the network topology used in this VPN example.

**Figure 43: Network Topology for the Interprovider VPN Example**



For configuration information see the following sections:

- [Configuration for Router A on page 320](#)
- [Configuration for Router B on page 321](#)
- [Configuration for Router C on page 322](#)
- [Configuration for Router D on page 323](#)
- [Configuration for Router E on page 324](#)
- [Configuration for Router F on page 326](#)

## Configuration for Router A

Configure a family `inet` EBGP session with Router B and export the direct routes:

```
[edit]
protocols {
  bgp {
    group to-provider {
      export attached;
      peer-as 10023;
      neighbor 192.168.198.2;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
```

```

    }
  }

```

## Configuration for Router B

Router A is configured as a CE router (using the **routing-instances** statement) in the configuration for Router B. Because they exchange VPN-IPv4 routes, Router D and Router C are configured as PE routers.

Configure Router B:

```

[edit]
protocols {
  rsvp {
    interface t3-0/0/0.0;
  }
  mpls {
    label-switched-path to-routerC {
      to 10.255.14.171;
      description "to-routerC for use with VPNs";
    }
    interface t3-0/0/0.0;
    interface so-1/2/0.0;
  }
  bgp {
    group to-ibgp {
      type internal;
      local-address 10.255.14.175;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.171;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface t3-0/0/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
}
routing-instances {
  vpna {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group to-ce {
          peer-as 9;

```

```

        neighbor 192.168.198.1;
    }
}
}
}
}
policy-options {
    policy-statement vpna-import {
        term 1 {
            from {
                protocol bgp;
                community vpna-comm;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    policy-statement vpna-export {
        term 1 {
            from protocol bgp;
            then {
                community add vpna-comm;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
    community vpna-comm members target:100:1001;
}

```

## Configuration for Router C

In the BGP protocol configuration for Router C, include the **keep all** statement. When this statement is included, BGP must store every route learned through BGP.

When you configure **keep all** or **keep none** and the peers support route refresh, the local speaker sends a refresh message and performs an import evaluation. For these peers, the sessions do not restart. To determine if a peer supports refresh, check for **Peer supports Refresh capability** in the output of the **show bgp neighbor** command.



**CAUTION:** If you configure **keep all** or **keep none** and the peer does not support session restart, the associated BGP sessions are restarted (flapped).

Configure two BGP sessions (configure **family inet-vpn** on both sessions):

- IBGP session to Router B (group **to-ibgp** in this example)
- EBGP session to Router D (group **to-ebgp-pe** in this example)

Interface **t3-0/2/0** is added at the **[edit protocols mpls]** hierarchy level, allowing BGP to announce routes with labels over the EBGP session.

Configure Router C:

```
[edit]
protocols {
  rsvp {
    interface t3-0/2/0.0;
  }
  mpls {
    label-switched-path to-routerB {
      to 10.255.14.175;
      description "to-routerB for use with vpns";
    }
    interface t3-0/2/0.0;
    interface so-0/0/0.0;
  }
  bgp {
    keep all;
    group to-ibgp {
      type internal;
      local-address 10.255.14.171;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.175;
    }
    group to-ebgp-pe {
      type external;
      family inet-vpn {
        unicast;
      }
      neighbor 192.168.197.22 {
        peer-as 10045;
      }
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface t3-0/2/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
}
```

## Configuration for Router D

The configuration for Router D is almost identical to that of Router C:

```
[edit]
protocols {
```

```
    rsvp {
      interface fe-1/1/0.0;
    }
    mpls {
      label-switched-path to-E {
        to 10.255.14.177;
        description "to-routerE for vpna";
      }
      interface fe-1/1/0.0;
      interface so-0/1/0.0;
    }
    bgp {
      keep all;
      group to-ibgp-pe {
        type internal;
        family inet-vpn {
          unicast;
        }
        neighbor 10.255.14.177;
      }
      group to-ebgp-pe {
        type external;
        family inet-vpn {
          unicast;
        }
        peer-as 10023;
        neighbor 192.168.197.21;
      }
    }
    ospf {
      traffic-engineering;
      reference-bandwidth 4g;
      area 0.0.0.0 {
        interface fe-1/1/0.0;
        interface lo0.0 {
          passive;
        }
      }
    }
  }
}
```

## Configuration for Router E

The configuration for Router E is very similar to the configuration for Router B:

```
[edit]
protocols {
  rsvp {
    interface fe-1/1/2.0;
  }
  mpls {
    label-switched-path to-routerD {
      to 10.255.14.173;
      description "to-routerD for use with VPNa";
    }
    interface fe-1/1/2.0;
  }
}
```



```

    interface so-1/2/0.0;
  }
  bgp {
    group to-ibgp-pe {
      type internal;
      local-address 10.255.14.177;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.173;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface fe-1/1/2.0;
      interface lo0.0 {
        passive;
      }
    }
  }
}
routing-instances {
  vpna {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 10.255.14.177:11;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group to-routerF-ce {
          neighbor 192.168.198.14 {
            peer-as 11;
          }
        }
      }
    }
  }
}
policy-options {
  policy-statement vpna-import {
    term 1 {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  policy-statement vpna-export {
    term 1 {

```

```

        from protocol bgp;
        then {
            community add vpna-comm;
            accept;
        }
    }
    term 2 {
        then reject;
    }
}
community vpna-comm members target:100:1001;
}

```

## Configuration for Router F

Configure Router F as a CE router; the configuration is similar to that for Router A:

```

[edit]
protocols {
    bgp {
        group to-provider {
            type external;
            export attached;
            neighbor 192.168.198.13 {
                peer-as 10045;
            }
        }
    }
}
policy-options {
    policy-statement attached {
        from protocol direct;
        then accept;
    }
}
}

```

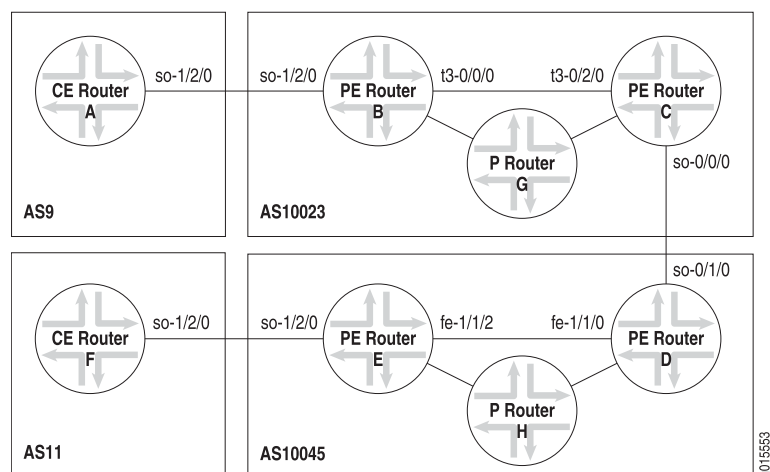
- Related Documentation**
- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
  - [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)

## Interprovider VPN Example—Multihop MP-EBGP with P Routers

In this example, labeled IPv4 (not VPN-IPv4), routes are exchanged by the AS border routers (Router C and Router D) to provide MPLS connectivity between the PE routers. Router G and H are provider routers.

[Figure 44 on page 327](#) illustrates the network topology used in this VPN example.

Figure 44: Network Topology of Interprovider VPN Example—Multihop MP-EBGP



Only routes internal to the service provider networks should be announced between Router C and Router D. Configure this by including the **family inet labeled-unicast** statement in the IBGP and EBGp configuration on the PE routers. When you set **family inet labeled-unicast**, the local router announces internal routes from inet.0 in the following manner:

- If a label exists for the route, the local router creates a label, performs a swap, and announces the route from inet.0 with the label.
- If a label does not exist for the route, the local router creates a label, performs a pop, and announces the route from inet.0 with the label.

Routes learned from the **labeled-unicast** session are placed into the inet.0 routing table.

In addition, you configure a multihop MP-EBGP session between the end PE routers (Router B and Router E). This additional MP-EBGP session allows the announcement of VPN-IPv4 routes, and allows you to maintain VPN connectivity while keeping VPN-IPv4 routes out of the core of the network.

For configuration information, see the following sections:

- [Configuration for Router A on page 327](#)
- [Configuration for Router B on page 328](#)
- [Configuration for Router C on page 329](#)
- [Configuration for Router D on page 331](#)
- [Configuration for Router E on page 332](#)
- [Configuration for Router F on page 333](#)

## Configuration for Router A

The configuration for Router A in this example is identical to the configuration for Router A in “Interprovider VPN Example—MP-EBGP Between ISP Peer Routers” on page 319. See “Configuration for Router A” on page 320

## Configuration for Router B

Router A is configured as a CE router (using the **routing-instances** statement) in the configuration for Router B. Because they exchange VPN-IPv4 routes, Router C and Router D are configured as PE routers.

In the BGP group **to-ibgp**, include the **family inet labeled-unicast** statement to pass labeled IPv4 routes, and configure an EBGP multihop session to pass VPN-IPv4 routes:

```
[edit]
protocols {
  bgp {
    group to-ibgp {
      type internal;
      local-address 10.255.14.175;
      family inet {
        labeled-unicast {
          resolve-vpn;
        }
      }
      neighbor 10.255.14.171;
    }
    group to-remote-pe {
      multihop {
        ttl 10;
      }
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.177 {
        peer-as 10045;
      }
    }
  }
  mpls {
    label-switched-path to-routerC {
      to 10.255.14.171;
      description "to-routerC for use with VPNs";
    }
    interface t3-0/0/0.0;
    interface so-1/2/0.0;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface t3-0/0/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
  rsvp {
    interface t3-0/0/0.0;
  }
}
```

```

routing-instances {
  vpna {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group to-ce {
          peer-as 9;
          neighbor 192.168.198.1;
        }
      }
    }
  }
}
policy-options {
  policy-statement vpna-import {
    term 1 {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  policy-statement vpna-export {
    term 1 {
      from protocol bgp;
      then {
        community add vpna-comm;
        accept;
      }
    }
    term 2 {
      then reject;
    }
  }
  community vpna-comm members target:100:1001;
}
}

```

## Configuration for Router C

Configure two BGP sessions (configure **family inet-vpn** on both sessions):

- IBGP session to Router B (group **to-ibgp** in this example)
- EBGP session to Router D (group **to-ebgp-pe** in this example)

Interface **t3-0/2/0** is added at the **[edit protocols mpls]** hierarchy level, allowing BGP to announce routes with labels over the EBGP session.

Configure Router C:

```
[edit]
protocols {
  bgp {
    group to-ibgp {
      type internal;
      local-address 10.255.14.171;
      family inet {
        labeled-unicast;
      }
      neighbor 10.255.14.175;
    }
    group to-ebgp-pe {
      type external;
      family inet {
        labeled-unicast;
      }
      export internal;
      neighbor 192.168.197.22 {
        peer-as 10045;
      }
    }
  }
  mpls {
    label-switched-path to-routerB {
      to 10.255.14.175;
      description "to-routerB for use with vpns";
    }
    interface t3-0/2/0.0;
    interface so-0/0/0.0;
    traffic-engineering bgp-igp;
  }
  rsvp {
    interface t3-0/2/0.0;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface t3-0/2/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
}
policy-options {
  policy-statement internal {
    term 1 {
      from protocol [ospf direct ldp];
      then accept;
    }
    term 2 {
      then reject;
    }
  }
}
```

```

    }
  }

```

## Configuration for Router D

Configure Router D:

```

[edit]
protocols {
  bgp {
    group to-ibgp-pe {
      type internal;
      family inet {
        labeled-unicast;
      }
      neighbor 10.255.14.177;
    }
    group to-ebgp-pe {
      type external;
      family inet {
        labeled-unicast;
      }
      export internal;
      peer-as 10023;
      neighbor 192.168.197.21;
    }
  }
  mpls {
    label-switched-path to-E {
      to 10.255.14.177;
      description "to-routerE for vpna";
    }
    interface fe-1/1/0.0;
    interface so-0/1/0.0;
    traffic-engineering bgp-igp;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface fe-1/1/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
  rsvp {
    interface fe-1/1/0.0;
  }
}
policy-options {
  policy-statement internal {
    term 1 {
      from protocol [ospf direct ldp];
      then accept;
    }
    term 2 {

```

```

        then reject;
    }
}
}
}

```

## Configuration for Router E

The configuration for Router E is very similar to the configuration for Router B:

```

[edit]
protocols {
  bgp {
    group to-ibgp-pe {
      type internal;
      local-address 10.255.14.177;
      family inet {
        labeled-unicast;
      }
      neighbor 10.255.14.173;
    }
    group to-remote-pe {
      multihop {
        ttl 10;
      }
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.175 {
        peer-as 10023;
      }
    }
  }
  mpls {
    label-switched-path to-routerD {
      to 10.255.14.173;
      description "to-routerD for use with VPNa";
    }
    interface fe-1/1/2.0;
    interface so-1/2/0.0;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface fe-1/1/2.0;
      interface lo0.0 {
        passive;
      }
    }
  }
  rsvp {
    interface fe-1/1/2.0;
  }
}
routing-instances {
  vpna {

```



```

instance-type vrf;
interface so-1/2/0.0;
route-distinguisher 10.255.14.177:11;
vrf-import vpna-import;
vrf-export vpna-export;
protocols {
  bgp {
    group to-routerF-ce {
      neighbor 192.168.198.14 {
        peer-as 11;
      }
    }
  }
}
}
policy-options {
  policy-statement vpna-import {
    term 1 {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  policy-statement vpna-export {
    term 1 {
      from protocol bgp;
      then {
        community add vpna-comm;
        accept;
      }
    }
    term 2 {
      then reject;
    }
  }
  community vpna-comm members target:100:1001;
}
}

```

## Configuration for Router F

The configuration for Router F in this example is identical to the configuration for Router F in [“Interprovider VPN Example—MP-EBGP Between ISP Peer Routers” on page 319](#). See [“Configuration for Router F” on page 326](#).

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)

## Example: Configuring Interprovider Layer 3 VPN Option A

---

Interprovider Layer 3 VPN Option A provides interprovider VRF-to-VRF connections at the AS boundary routers (ASBRs). Compared to Option B and Option C, Option A is the least scalable solution.

This example provides a step-by-step procedure to configure interprovider Layer 3 VPN option A, which is one of the recommended implementations of MPLS VPN when that service is required by a customer that has more than one AS and but not all of the customer's ASs can be serviced by the same service provider. It is organized in the following sections:

- [Requirements on page 334](#)
- [Overview and Topology on page 334](#)
- [Configuration on page 335](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later.
- Eight M Series, T Series, TX Series, or MX Series Juniper Networks routers.

### Overview and Topology

This is the simplest and least scalable interprovider VPN solution to the problem of providing VPN services to a customer that has different sites, not all of which can use the same service provider (SP).

*RFC 4364*, section 10, refers to this method as Interprovider VRF-to-VRF connections at the AS border routers.

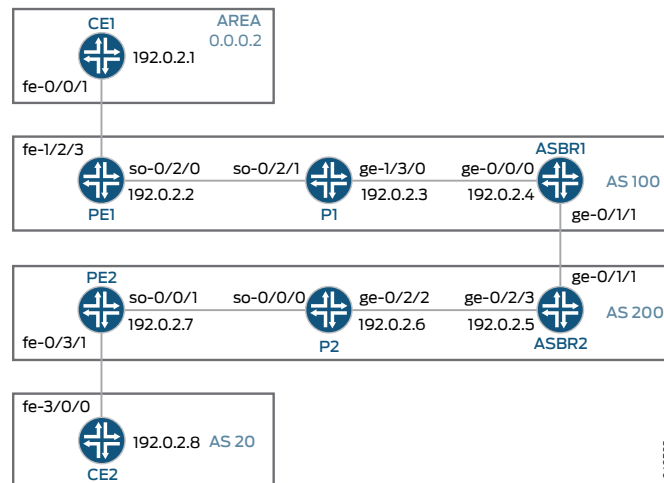
In this configuration:

- The virtual routing and forwarding (VRF) table in the ASBR of one AS is linked to the VRF table in the ASBR in the other AS. Each ASBR must contain a VRF instance for every VPN configured in both service provider networks. Then an IGP or BGP must be configured between the ASBRs. This has the disadvantage of limiting scalability.
- In this configuration, the autonomous system boundary routers (ASBRs) at both SPs are configured as regular PE routers, and provide MPLS L3 VPN service to the neighbor SP.
- Each PE router treats the other as if it were a customer edge (CE) router. ASBRs play the role of regular CE routers for the ASBR of the remote SP. ASBRs see each other as CE devices.
- A provider edge (PE) router in one autonomous system (AS) attaches directly to a PE router in another AS.

- The two PE routers are attached by multiple sub-interfaces, at least one for each of the VPNs whose routes need to be passed from AS to AS.
- The PE routers associate each sub-interface with a VPN routing and forwarding (VRF) table, and use EBGP to distribute unlabeled IPv4 addresses to each other.
- In this solution, all common VPNs defined at both PEs must also be defined at one or more ASBRs between the two SPs. This is not a very scalable methodology, especially when a transit SP is used by two regional SPs for interconnection.
- This is a procedure that is simple to configure and it does not require MPLS at the border between ASs. Additionally, it does not scale as well as other recommended procedures.

The topology of the network is shown in [Figure 45 on page 335](#).

*Figure 45: Physical Topology of Interprovider Layer 3 VPN Option A*



## Configuration



**NOTE:** The procedure presented here is written with the assumption that the reader is already familiar with MPLS MVPN configuration. This example focuses on explaining the unique configuration required for carrier-of-carriers solutions for VPN services to different sites.

To configure interprovider layer 3 VPN option A, perform the following tasks:

- [Configuring Router CE1 on page 336](#)
- [Configuring Router PE1 on page 336](#)
- [Configuring Router P1 on page 339](#)
- [Configuring Router ASBR1 on page 340](#)
- [Configuring Router ASBR2 on page 342](#)

- [Configuring Router P2 on page 344](#)
- [Configuring Router PE2 on page 345](#)
- [Configuring Router CE2 on page 348](#)
- [Verifying the VPN Operation on page 349](#)

---

### Configuring Router CE1

---

#### Step-by-Step Procedure

1. On Router CE1, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE1 and Router PE1. Specify the **inet** address family type.  

```
[edit interfaces fe-0/0/1.0]
family inet {
  address 198.51.100.1/24;
}
```
2. On Router CE1, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.  

```
[edit interfaces lo0]
unit 0 {
  family inet {
    address 192.0.2.1/32;
  }
}
```
3. On Router CE1, configure a routing protocol. The routing protocol can be a static route, RIP, OSPF, ISIS, or EBGp. In this example we configure OSPF. Include the Fast Ethernet interface for the link between Router CE1 and Router PE1 and the logical loopback interface of Router CE1.  

```
[edit protocols]
ospf {
  area 0.0.0.2 {
    interface fe-0/0/1.0;
    interface lo0.0;
  }
}
```

---

### Configuring Router PE1

---

#### Step-by-Step Procedure

1. On Router PE1, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET and Fast Ethernet interfaces.  

```
[edit interfaces]
so-0/2/0 {
  unit 0 {
    family inet {
      address 192.168.1.9/24;
    }
  }
}
```

```

        family mpls;
    }
}
fe-1/2/3 {
    unit 0 {
        family inet {
            address 198.51.100.2/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.2/32;
        }
    }
}
}

```

2. On Router PE1, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the OSPF protocol within the VRF. Specify the customer-facing Fast Ethernet interface and specify the export policy to export BGP routes into OSPF.

```

[edit routing-instances]
vpn2CE1 {
    instance-type vrf;
    interface fe-1/2/3.0;
    route-distinguisher 1:100;
    vrf-import vpnimport;
    vrf-export vpnexport;
    protocols {
        ospf {
            export bgp-to-ospf;
            area 0.0.0.2 {
                interface fe-1/2/3.0;
            }
        }
    }
}
}

```

3. On Router PE1, configure the RSVP and MPLS protocols to support the label-switched path (LSP). Configure the LSP to Router ASBR1 and specify the IP address of the logical loopback interface on Router ASBR1. Configure a BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE1. Specify the neighbor address as the logical loopback interface on Router ASBR1. Specify the **inet-vpn** address family and **unicast** traffic type to enable BGP to carry IPv4 network layer reachability information (NLRI) for VPN routes. Configure the OSPF protocol. Specify the core-facing SONET interface and specify the logical loopback interface on Router PE1.

```
[edit protocols]
rsvp {
  interface so-0/2/0.0;
  interface lo0.0;
}
mpls {
  label-switched-path To-ASBR1 {
    to 192.0.2.4;
  }
  interface so-0/2/0.0;
  interface lo0.0;
}
bgp {
  group To_ASBR1 {
    type internal;
    local-address 192.0.2.2;
    neighbor 192.0.2.4 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface so-0/2/0.0;
    interface lo0.0;
  }
}
```

4. On Router PE1, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 100;
```

5. On Router PE1, configure a policy to export the BGP routes into OSPF.

```
[edit policy-options]
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

6. On Router PE1, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```
[edit policy-options]
policy-statement vpnexport {
  term 1 {
```

```

    from protocol ospf;
    then {
        community add test_comm;
        accept;
    }
}
term 2 {
    then reject;
}
}

```

7. On Router PE1, configure a policy to import routes from BGP that have the **test\_comm** community attached.

```

[edit policy-options]
policy-statement vpnimport {
    term 1 {
        from {
            protocol bgp;
            community test_comm;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}

```

8. On Router PE1, define the **test\_comm** BGP community with a route target.

```

[edit policy-options]
community test_comm members target:1:100;

```

### Configuring Router P1

#### Step-by-Step Procedure

1. On Router P1, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
so-0/2/1 {
    unit 0 {
        family inet {
            address 192.168.1.4/24;
        }
        family mpls;
    }
}
ge-1/3/0 {
    unit 0 {
        family inet {
            address 192.168.2.5/24;
        }
    }
}

```

```

    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.3/32;
    }
  }
}
}

```

2. On Router P1, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
  interface so-0/2/1.0;
  interface ge-1/3/0.0;
  interface lo0.0;
}
mpls {
  interface lo0.0;
  interface ge-1/3/0.0;
  interface so-0/2/1.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-1/3/0.0;
    interface so-0/2/1.0;
    interface lo0.0;
  }
}
}

```

## Configuring Router ASBR1

### Step-by-Step Procedure

1. On Router ASBR1, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** addresses families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.2.6/24;
    }
    family mpls;
  }
}

```



```

}
ge-0/1/1 {
  unit 0 {
    family inet {
      address 192.168.3.7/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.4/32;
    }
  }
}
}

```

2. On Router ASBR1, configure the **To\_ASBR2** routing instance. Specify the **vrf** instance type and specify the core-facing Gigabit Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Configure a route target for the VPN. Configure the BGP peer group within the VRF. Specify AS 200 as the peer AS and specify the IP address of the Gigabit Ethernet interface on Router ASBR2 as the neighbor address.

```

[edit routing instances]
To_ASBR2{
  instance-type vrf;
  interface ge-0/1/1.0;
  route-distinguisher 1:100;
  vrf-target target:1:100;
  protocols {
    bgp {
      group To_ASBR2 {
        type external;
        neighbor 192.168.3.8 {
          peer-as 200;
        }
      }
    }
  }
}

```

3. On Router ASBR1, configure the RSVP and MPLS protocols to support the LSP by specifying the Gigabit Ethernet interface that is facing the P1 router.

Configure the OSPF protocol by specifying the Gigabit Ethernet interface that is facing the P1 router and the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
  interface ge-0/0/0.0;
  interface lo0.0;
}

```

```

mpls {
  label-switched-path To-PE1 {
    to 192.0.2.2;
  }
  interface lo0.0;
  interface ge-0/0/0.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
  }
}

```

4. On Router ASBR1, create the **To-PE1** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE1.

```

[edit protocols]
bgp {
  group To-PE1 {
    type internal;
    local-address 192.0.2.4;
    neighbor 192.0.2.2 {
      family inet-vpn {
        unicast;
      }
    }
  }
}

```

5. On Router ASBR1, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 100;

```

## Configuring Router ASBR2

### Step-by-Step Procedure

1. On Router ASBR2, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
ge-0/1/1 {
  unit 0 {
    family inet {
      address 192.168.3.8/24;
    }
    family mpls;
  }
}

```

```

ge-0/2/3 {
  unit 0 {
    family inet {
      address 192.168.4.10/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.5/32;
    }
  }
}

```

2. On Router ASBR2, configure the **To\_ASBR1** routing instance. Specify the **vrf** instance type and specify the core-facing Gigabit Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Configure a route target for the VPN. Configure the BGP peer group within the VRF. Specify AS 100 as the peer AS and specify the IP address of the Gigabit Ethernet interface on Router ASBR1 as the neighbor address.

```

[edit routing-instances]
To_ASBR1 {
  instance-type vrf;
  interface ge-0/1/1.0;
  route-distinguisher 1:100;
  vrf-target target:1:100;
  protocols {
    bgp {
      group To_ASBR1 {
        type external;
        neighbor 192.168.3.7 {
          peer-as 100;
        }
      }
    }
  }
}

```

3. On Router ASBR2, configure the RSVP and MPLS protocols to support the LSP by specifying the Gigabit Ethernet interface that is facing the P2 router.

Configure the OSPF protocol by specifying the Gigabit Ethernet interface that is facing the P2 router and the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
  interface ge-0/2/3.0;
  interface lo0.0;
}
mpls {

```

```

label-switched-path To_PE2 {
  to 192.0.2.7;
}
interface lo0.0;
interface ge-0/2/3.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/2/3.0;
    interface lo0.0;
  }
}

```

4. On Router ASBR2, create the **To-PE2** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE2.

```

[edit protocols]
bgp {
  group To-PE2 {
    type internal;
    local-address 192.0.2.5;
    neighbor 192.0.2.7 {
      family inet-vpn {
        unicast;
      }
    }
  }
}

```

5. On Router ASBR2, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 200;

```

## Configuring Router P2

### Step-by-Step Procedure

1. On Router P2, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.5.11/24;
    }
    family mpls;
  }
}
ge-0/2/2 {
  unit 0 {

```

```

        family inet {
            address 192.168.4.12/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.6/32;
        }
    }
}
}

```

2. On Router P2, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
    interface so-0/0/0.0;
    interface ge-0/2/2.0;
    interface lo0.0;
}
mpls {
    interface lo0.0;
    interface ge-0/2/2.0;
    interface so-0/0/0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/2/2.0;
        interface so-0/0/0.0;
        interface lo0.0;
    }
}
}

```

## Configuring Router PE2

### Step-by-Step Procedure

1. On Router PE2, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET and Fast Ethernet interfaces.

```

[edit interfaces]
so-0/0/1 {
    unit 0 {
        family inet {
            address 192.168.5.12/24;
        }
        family mpls;
    }
}

```

```

    }
  }
  fe-0/3/1 {
    unit 0 {
      family inet {
        address 192.168.6.13/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.7/32;
      }
    }
  }
}

```

2. On Router PE2, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the BGP peer group within the VRF. Specify AS **20** as the peer AS and specify the IP address of the Fast Ethernet interface on Router CE2 as the neighbor address.

```

[edit routing-instances]
vpn2CE2 {
  instance-type vrf;
  interface fe-0/3/1.0;
  route-distinguisher 1:100;
  vrf-import vpnimport;
  vrf-export vpnexport;
  protocols {
    bgp {
      group To_CE2 {
        peer-as 20;
        neighbor 192.168.6.14;
      }
    }
  }
}

```

3. On Router PE2, configure the RSVP and MPLS protocols to support the LSP. Configure the LSP to ASBR2 and specify the IP address of the logical loopback interface on Router ASBR2. Configure a BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE2. Specify the neighbor address as the logical loopback interface on the Router ASBR2. Specify the **inet-vpn** address family and **unicast** traffic type to enable BGP to carry IPv4 NLRI for VPN routes. Configure the OSPF protocol. Specify the core-facing SONET interface and specify the logical loopback interface on Router PE2.

```

[edit protocols]
rsvp {
  interface so-0/0/1.0;
}

```

```

    interface lo0.0;
  }
  mpls {
    label-switched-path To-ASBR2 {
      to 192.0.2.5;
    }
    interface so-0/0/1.0;
    interface lo0.0;
  }
  bgp {
    group To_ASBR2 {
      type internal;
      local-address 192.0.2.7;
      neighbor 192.0.2.5 {
        family inet-vpn {
          unicast;
        }
      }
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/0/1.0;
      interface lo0.0;
    }
  }
}

```

4. On Router PE2, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 200;

```

5. On Router PE2, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```

[edit policy-options]
policy-statement vpnexport {
  term 1 {
    from protocol bgp;
    then {
      community add test_comm;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}

```

6. On Router PE2, configure a policy to import routes from BGP that have the **test\_comm** community attached.

```

[edit policy-options]
policy-statement vpnimport {

```

```

term 1 {
  from {
    protocol bgp;
    community test_comm;
  }
  then accept;
}
term 2 {
  then reject;
}
}

```

7. On Router PE2, define the **test\_comm** BGP community with a route target.

```

[edit policy-options]
community test_comm members target:1:100;

```

### Configuring Router CE2

#### Step-by-Step Procedure

1. On Router CE2, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE2 and Router PE2. Specify the **inet** address family type.

```

[edit interfaces]
fe-3/0/0 {
  unit 0 {
    family inet {
      address 192.168.6.14/24;
    }
  }
}

```

2. On Router CE2, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.

```

[edit interfaces lo0]
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.8/32;
    }
  }
}

```

3. On Router CE2, define a policy named **myroutes** that accepts direct routes.

```

[edit policy-options]
policy-statement myroutes {
  from protocol direct;
  then accept;
}

```



- On Router CE2, configure a routing protocol. The routing protocol can be a static route, RIP, OSPF, ISIS, or EBGp. In this example, we configure EBGp. Specify AS 200 as the peer AS and specify the BGP neighbor IP address as the Fast Ethernet interface of Router PE2.

```
[edit protocols]
bgp {
  group To_PE2 {
    neighbor 192.168.6.13 {
      export myroutes;
      peer-as 200;
    }
  }
}
```

- On Router CE2, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 20;
```

### Verifying the VPN Operation

#### Step-by-Step Procedure

- Commit the configuration on each router.



**NOTE:** The MPLS labels shown in this example will be different than the labels used in your configuration.

- On Router PE1, display the routes for the **vpn2CE1** routing instance using the **show ospf route** command. Verify that the **192.0.2.1** route is learned from OSPF.

```
user@PE1> show ospf route instance vpn2CE1
Topology default Route Table:
```

Prefix	Path	Route	NH	Metric	NextHop	NextHop
	Type	Type	Type		Interface	addr/label
<b>192.0.2.1</b>	Intra	Router	IP	1	fe-1/2/3.0	198.51.100.1
192.0.2.1/32	Intra	Network	IP		1 fe-1/2/3.0	
198.51.100.1						
198.51.100.0/24	Intra	Network	IP	1	fe-1/2/3.0	
198.51.100.1						

- On Router PE1, use the **show route advertising-protocol** command to verify that Router PE1 advertises the **192.0.2.1** route to Router ASBR1 using MP-BGP with the VPN MPLS label.

```
user@PE1> show route advertising-protocol bgp 192.0.2.4 extensive
vpn2CE1.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group To_PE1 type Internal
```

```

Route Distinguisher: 1:100
VPN Label: 299856
Nexthop: Self
Flags: Nexthop Change
MED: 1
Localpref: 100
AS path: [100] I
Communities: target:1:100 rte-type:0.0.0.2:1:0

```

4. On Router ASBR1, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.1** route and places it in the **To\_ASBR2.inet.0** routing table.

```

user@ASBR1> show route receive-protocol bgp 192.0.2.2 extensive
inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

To_ASBR2.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  Route Distinguisher: 1:100
  VPN Label: 299856
  Nexthop: 192.0.2.2
  MED: 1
  Localpref: 100
  AS path: I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

MPLS.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

BGP.13VPN.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

* 1:100:192.0.2.1/32 (1 entry, 0 announced)
  Route Distinguisher: 1:100
  VPN Label: 299856
  Nexthop: 192.0.2.2
  MED: 1
  Localpref: 100
  AS path: I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

5. On Router ASBR1, use the **show route advertising-protocol** command to verify that Router ASBR1 advertises the **192.0.2.1** route to Router ASBR2.

```

user@ASBR1> show route advertising-protocol bgp 192.168.3.8 extensive
To_ASBR2.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group To_ASBR2.inet.0 type External
  Nexthop: Self
  AS path: [100] I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

6. On Router ASBR2, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.1** route and places it in the **To\_ASBR1.inet.0** routing table.

```

user@ASBR2> show route receive-protocol bgp 192.168.3.7 extensive

```

```

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

To_ASBR1.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  Accepted
  Nexthop: 192.168.3.7
  AS path: 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

MPLS.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

BGP.13VPN.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

7. On Router ASBR2, use the **show route advertising-protocol** command to verify that Router ASBR2 advertises the **192.0.2.1** route to Router PE2.

```

user@ASBR2> show route advertising-protocol bgp 192.0.2.7 extensive
To_ASBR1.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group To-PE2 type Internal
  Route Distinguisher: 1:100
  VPN Label: 299936
  Nexthop: Self
  Flags: Nexthop Change
  Localpref: 100
  AS path: [200] 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

8. On Router PE2, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.1** route and places it in the **vpn2CE2.inet.0** routing table.

```

user@PE2> show route receive-protocol bgp 192.0.2.5 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

__juniper_private1__.inet.0: 14 destinations, 14 routes (8 active, 0 holddown,
6 hidden)

__juniper_private2__.inet.0: 1 destinations, 1 routes (0 active, 0 holddown,
1 hidden)

vpn2CE2.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 299936
  Nexthop: 192.0.2.5
  Localpref: 100
  AS path: 100 I
  AS path: Recorded
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

9. On Router PE2, use the **show route advertising-protocol** command to verify that Router PE2 advertises the **192.0.2.1** route to Router CE2 through the **To\_CE2** peer group.

```
user@PE2> show route advertising-protocol bgp 192.168.6.14 extensive
vpn2CE2.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group To_CE2 type External
    Nexthop: Self
    AS path: [200] 100 I
    Communities: target:1:100 rte-type:0.0.0.2:1:0
```

10. On Router CE2, use the **show route** command to verify that Router CE2 receives the **192.0.2.1** route from Router PE2.

```
user@CE2> show route 192.0.2.1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32          *[BGP/170] 00:25:36, localpref 100
                     AS path: 200 100 I
                     > to 192.168.6.13 via fe-3/0/0.0
```

11. On Router CE2, use the **ping** command and specify **192.0.2.8** as the source of the ping packets to verify connectivity with Router CE1.

```
user@CE2> ping 192.0.2.1 source 192.0.2.8
PING 192.0.2.1 (192.0.2.1): 56 data bytes
64 bytes from 192.0.2.1: icmp_seq=0 ttl=58 time=4.672 ms
64 bytes from 192.0.2.1: icmp_seq=1 ttl=58 time=10.480 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=58 time=10.560 ms
```

12. On Router PE2, use the **show route** command to verify that the traffic is sent with an inner label of **299936** and a top label of **299776**.

```
user@PE2> show route 192.0.2.1 detail

vpn2CE2.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
192.0.2.1/32 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 1:100
    Next hop type: Indirect
    Next-hop reference count: 6
    Source: 192.0.2.5
    Next hop type: Router, Next hop index: 648
    Next hop: via so-0/0/1.0 weight 0x1, selected
    Label-switched-path To-ASBR2
    Label operation: Push 299936, Push 299776(top)
    Protocol next hop: 192.0.2.5
    Push 299984
    Indirect next hop: 8c6109c 262143
    State: <Secondary Active Int Ext>
    Local AS: 200 Peer AS: 200
    Age: 3:37 Metric2: 2
    Task: BGP_200.192.0.2.5+179
    Announcement bits (3): 0-RT 1-KRT 2-BGP RT Background
    AS path: 100 I
    AS path: Recorded
```

```

Communities: target:1:100 rte-type:0.0.0.2:1:0
Accepted
VPN Label: 299984
Localpref: 100
Router ID: 192.0.2.5
Primary Routing Table BGP.13VPN.0

```

13. On Router ASBR2, use the **show route table** command to verify that Router ASBR2 receives the traffic.

```

user@ASBR2# show route table mpls.0 detail
299936 (1 entry, 1 announced)
  *VPN      Preference: 170
            Next hop type: Router, Next hop index: 649
            Next-hop reference count: 2
            Source: 192.168.3.7           Next hop: 192.168.3.7 via
ge-0/1/1.0, selected
            Label operation: Pop
            State: <Active Int Ext>
            Local AS: 200
            Age: 9:54
            Task: BGP RT Background
            Announcement bits (1): 0-KRT
            AS path: 100 I
            Ref Cnt: 1
            Communities: target:1:100 rte-type:0.0.0.2:1:0

```

14. On Router ASBR2, use the **show route table** command to verify that Router ASBR2 receives the traffic.

```

user@ASBR2# show route 192.0.2.1 detail
To_ASBR1.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
192.0.2.1/32 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Next hop type: Router, Next hop index: 576
            Next-hop reference count: 3
            Source: 192.168.3.7
            Next hop: 192.168.3.7 via ge-0/1/1.0, selected
            State: <Active Ext>
            Peer AS: 100
            Age: 13:07
            Task: BGP_192.168.3.7+53372
            Announcement bits (2): 0-KRT 1-BGP RT Background
            AS path: 100 I
            Communities: target:1:100 rte-type:0.0.0.2:1:0
            Accepted
            Localpref: 100
            Router ID: 192.168.3.7

```

15. On Router ASBR1, use the **show route** command to verify that ASBR1 sends traffic toward PE1 with the top label **299792** and VPN label **299856**.

```

user@ASBR1# show route 192.0.2.1 detail
To_ASBR2.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
192.0.2.1/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 1:100
            Next hop type: Indirect

```

```

Next-hop reference count: 3
Source: 192.0.2.2
Next hop type: Router, Next hop index: 669
Next hop: 192.168.2.5 via ge-0/0/0.0 weight 0x1, selected
Label-switched-path To_PE1
Label operation: Push 299856, Push 299792(top)
Protocol next hop: 192.0.2.2                      Push 299856
Indirect next hop: 8af70a0 262143
State: <Secondary Active Int Ext>
Local AS: 100 Peer AS: 100
Age: 12:15 Metric: 1 Metric2: 2
Task: BGP_100.192.0.2.2+58065
Announcement bits (2): 0-KRT 1-BGP RT Background
AS path: I
Communities: target:1:100 rte-type:0.0.0.2:1:0
VPN Label: 299856
Localpref: 100
Router ID: 192.0.2.2
Primary Routing Table BGP.13VPN.0

```

16. On Router PE1, use the **show route table** command to verify that Router PE1 receives the traffic with label **299856**, pops the label, and the traffic is sent toward Router CE1 through interface **fe-1/2/3.0**.

```

lab@PE1# show route table mpls.0 detail
299856 (1 entry, 1 announced)
  *VPN Preference: 170
    Next hop type: Router, Next hop index: 666
    Next-hop reference count: 2
    Next hop: 198.51.100.8 via fe-1/2/3.0, selected
    Label operation: Pop
    State: <Active Int Ext>
    Local AS: 100
    Age: 17:38
    Task: BGP RT Background
    Announcement bits (1): 0-KRT
    AS path: I
    Ref Cnt: 1
    Communities: rte-type:0.0.0.2:1:0

```

17. On Router PE1, use the **show route** command to verify that PE1 receives the traffic after the top label is popped by Router P and the traffic is sent toward Router CE1 through interface **fe-1/2/3.0**.

```

lab@PE1# show route 192.0.2.1 detail

vpn2CE1.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
192.0.2.1/32 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next hop type: Router, Next hop index: 634
    Next-hop reference count: 3
    Next hop: 198.51.100.8 via fe-1/2/3.0, selected
    State: <Active Int>
    Age: 18:42 Metric: 1
    Area: 0.0.0.2
    Task: VPN2alice-OSPFv2
    Announcement bits (2): 2-KRT 3-BGP RT Background
    AS path: I
    Communities: rte-type:0.0.0.2:1:0

```

## Related Documentation

### Example: Configuring Interprovider Layer 3 VPN Option B

Interprovider Layer 3 VPN Option B provides interprovider EBGP redistribution of labeled VPN-IPv4 routes from AS to neighboring AS. This solution is considered to be more scalable than Option A, but not as scalable as Option C.

This example provides a step-by-step procedure to configure interprovider layer 3 VPN option B, which is one of the recommended implementations of an MPLS VPN when that service is required by a customer that has more than one AS, but not all of the customer's ASs can be serviced by the same service provider. It is organized in the following sections:

- [Requirements on page 355](#)
- [Configuration Overview and Topology on page 355](#)
- [Configuration on page 357](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later.
- Eight M Series, T Series, TX Series, or MX Series Juniper Networks routers.

## Configuration Overview and Topology

Interprovider layer 3 VPN option B is a somewhat scalable solution to the problem of providing VPN services to a customer that has different sites, not all of which can use the same service provider. *RFC 4364*, section 10, refers to this method as interprovider EBGP redistribution of labeled VPN-IPv4 routes from AS to neighboring AS.

In the topology shown in Figure 1, the following events occur:

- The PE routers use IBGP to redistribute labeled VPN-IPv4 routes either to an ASBR, or to a route reflector of which an ASBR is a client.
- The ASBR then uses EBGP to redistribute those labeled VPN-IPv4 routes to an ASBR in another AS, which distributes them to the PE routers in that AS, or to another ASBR for distribution.
- Labeled VPN-IPv4 routes are distributed between ASBR routers on each site. There is no need to define a separate VPN routing and forwarding instance (VRF) for each common VPN that resides on two different SPs.
- Router PE2 distributes VPN-IPv4 routes to Router ASBR2 using MP-IBGP.
- Router ASBR2 distributes these labeled VPN-IPv4 routes to Router ASBR1, using the MP-EBGP session between them.
- Router ASBR1 redistributes those routes to Router PE1, using MP-IBGP. Each time a label is advertised, routers change the next-hop information and labels.

- An MPLS path is established between Router PE1 and Router PE2. This path enables changing of the next-hop attribute for the routes that are learned from the neighbor SP router and map the incoming label for the given routes to the outgoing label advertised to PE routers in the internal network.
- The ingress PE router inserts two labels onto the IP packet coming from the end customer. The inner label is for the VPN-IPv4 routes learned from internal ASBRs and the outer label is for the route to the internal ASBR, obtained through resource reservation protocol (RSVP) or label distribution protocol (LDP).
- When a packet arrives at the ASBR, it removes the outer label (when explicit-null signaling is used; otherwise, penultimate hop-popping (PHP) pops the label) and swaps the inner label with the label obtained from the neighbor ASBR through MP-EBGP label and prefix advertisements.
- The second ASBR swaps the VPN-IPv4 label and pushes another label to reach the PE router in its own AS.
- The remaining process is the same as for a regular VPN.

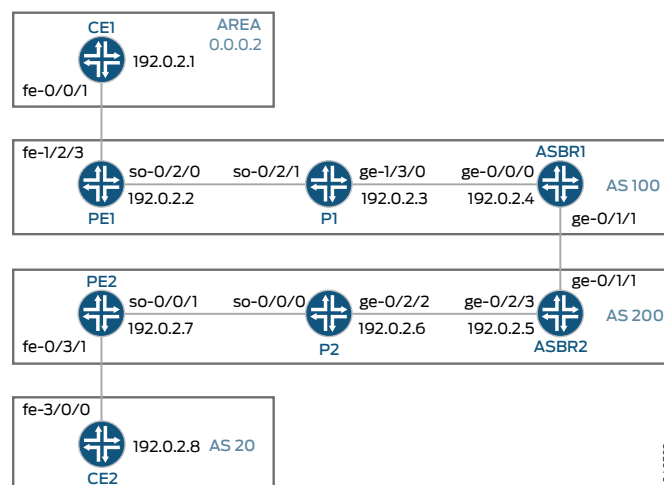


**NOTE:** In this solution, ASBR routers keep all VPN-IPv4 routes in the routing information base (RIB), and the labels associated with the prefixes are kept in the forwarding information base (FIB). Because the RIB and FIB tables can take occupy much of the respective allocated memory, this solution is not very scalable for an interprovider VPN.

If a transit SP is used between SP1 and SP2, the transit SP also has to keep all VPN-IPv4 routes in the RIB and the corresponding labels in the FIB. The ASBRs at the transit SP have the same functionality as ASBRs in the SP1 or SP2 networks in this solution.

The topology of the network is shown in [Figure 46 on page 356](#).

**Figure 46: Physical Topology of Interprovider Layer 3 VPN Option B**





## Configuration



**NOTE:** The procedure presented here is written with the assumption that the reader is already familiar with MPLS MVPN configuration. This example focuses on explaining the unique configuration required for carrier-of-carriers solutions for VPN services to different sites.

To configure layer 3 VPN option B, perform the following tasks:

- [Configuring Router CE1 on page 357](#)
- [Configuring Router PE1 on page 358](#)
- [Configuring Router P1 on page 361](#)
- [Configuring Router ASBR1 on page 362](#)
- [Configuring Router ASBR2 on page 364](#)
- [Configuring Router P2 on page 365](#)
- [Configuring Router PE2 on page 367](#)
- [Configuring Router CE2 on page 369](#)
- [Verifying the VPN Operation on page 370](#)

### Configuring Router CE1

#### Step-by-Step Procedure

1. On Router CE1, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE1 and Router PE1. Specify the **inet** address family type.
 

```
[edit interfaces fe-0/0/1.0]
family inet {
  address 203.0.113.1/24;
}
```
2. On Router CE1, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.
 

```
[edit interfaces lo0]
unit 0 {
  family inet {
    address 192.0.2.1/32;
  }
}
```
3. On Router CE1, configure a routing protocol. Include the logical interface for the link between Router CE1 and Router PE1 and the logical loopback interface of Router CE1. The routing protocol can be a static route, RIP, OSPF, ISIS, or EBGp. In this example we configure OSPF.
 

```
[edit protocols]
ospf {
```

```

        area 0.0.0.2 {
            interface fe-0/0/1.0;
            interface lo0.0;
        }
    }
}

```

## Configuring Router PE1

### Step-by-Step Procedure

1. On Router PE1, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET and Fast Ethernet interfaces.

```

[edit interfaces]
so-0/2/0 {
    unit 0 {
        family inet {
            address 192.168.1.2/24;
        }
        family mpls;
    }
}
fe-1/2/3 {
    unit 0 {
        family inet {
            address 203.0.113.3/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.2/32;
        }
    }
}
}

```

2. On Router PE1, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the OSPF protocol within the VRF. Specify the customer-facing Fast Ethernet interface and specify the export policy to export BGP routes into OSPF.

```

[edit routing-instances]
vpn2CE1 {
    instance-type vrf;
    interface fe-1/2/3.0;
    route-distinguisher 1:100;
    vrf-import vpnimport;
    vrf-export vpnexport;
    protocols {
        ospf {

```

```

        export bgp-to-ospf;
        area 0.0.0.2 {
            interface fe-1/2/3.0;
        }
    }
}

```

3. On Router PE1, configure the RSVP and MPLS protocols to support the label-switched path (LSP). Configure the LSP to Router ASBR1 and specify the IP address of the logical loopback interface on Router ASBR1. Configure a BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE1. Specify the neighbor address as the logical loopback interface on Router ASBR1. Specify the **inet-vpn** address family and **unicast** traffic type to enable BGP to carry IPv4 network layer reachability information (NLRI) for VPN routes. Configure the OSPF protocol. Specify the core-facing SONET interface and specify the logical loopback interface on Router PE1.

```

[edit protocols]
rsvp {
    interface so-0/2/0.0;
    interface lo0.0;
}
mpls {
    label-switched-path To-ASBR1 {
        to 192.0.2.4;
    }
    interface so-0/2/0.0;
    interface lo0.0;
}
bgp {
    group To_ASBR1 {
        type internal;
        local-address 192.0.2.2;
        neighbor 192.0.2.4 {
            family inet-vpn {
                unicast;
            }
        }
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface so-0/2/0.0;
        interface lo0.0;
    }
}

```

4. On Router PE1, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 100;

```

5. On Router PE1, configure a policy to export the BGP routes into OSPF.

```
[edit policy-options]
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

6. On Router PE1, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```
[edit policy-options]
policy-statement vpnexport {
  term 1 {
    from protocol ospf;
    then {
      community add test_comm;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}
```

7. On Router PE1, configure a policy to import routes from BGP that have the **test\_comm** community attached.

```
[edit policy-options]
policy-statement vpnimport {
  term 1 {
    from {
      protocol bgp;
      community test_comm;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

8. On Router PE1, define the **test\_comm** BGP community with a route target.

```
[edit policy-options]
community test_comm members target:1:100;
```

## Configuring Router P1

### Step-by-Step Procedure

1. On Router P1, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
so-0/2/1 {
  unit 0 {
    family inet {
      address 192.168.1.4/24;
    }
    family mpls;
  }
}
ge-1/3/0 {
  unit 0 {
    family inet {
      address 192.168.2.5/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.3/32;
    }
  }
}
```

2. On Router P1, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```
[edit protocols]
rsvp {
  interface so-0/2/1.0;
  interface ge-1/3/0.0;
  interface lo0.0;
}
mpls {
  interface lo0.0;
  interface ge-1/3/0.0;
  interface so-0/2/1.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-1/3/0.0;
```

```

        interface so-0/2/1.0;
        interface lo0.0;
    }
}

```

### Configuring Router ASBR1

- Step-by-Step Procedure**
1. On Router ASBR1, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
ge-0/0/0 {
    unit 0 {
        family inet {
            address 192.168.2.7/24;
        }
        family mpls;
    }
}
ge-0/1/1 {
    unit 0 {
        family inet {
            address 192.168.3.8/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.4/32;
        }
    }
}

```

2. On Router ASBR1, configure the RSVP and MPLS protocols to support the LSP by specifying the Gigabit Ethernet interface facing the P1 router and the **lo0.0** logical loopback interface.

Configure the OSPF protocol by specifying the Gigabit Ethernet interface that is facing the P1 router and the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
    interface ge-0/0/0.0;
    interface lo0.0;
}
mpls {
    label-switched-path To_PE1 {
        to 192.0.2.2;
    }
}

```

```

    }
    interface lo0.0;
    interface ge-0/0/0.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface ge-0/0/0.0;
      interface lo0.0;
    }
  }
}

```

3. On Router ASBR1, create the **To-PE1** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE1.

```

[edit protocols]
bgp {
  group To-PE1 {
    type internal;
    local-address 192.0.2.4;
    neighbor 192.0.2.2 {
      family inet-vpn {
        unicast;
      }
    }
  }
}

```

4. On Router ASBR1, create the **To-ASBR2** external BGP peer group. Enable the router to use BGP to advertise NLRI for unicast routes. Specify the neighbor IP peer address as the Gigabit Ethernet interface address of Router ASBR2.

```

[edit protocols]
bgp {
  group To-ASBR2 {
    type external;
    family inet-vpn {
      unicast;
    }
    neighbor 192.168.3.6 {
      peer-as 200;
    }
  }
}

```

5. On Router ASBR1, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 100;

```

## Configuring Router ASBR2

### Step-by-Step Procedure

1. On Router ASBR2, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
ge-0/1/1 {
  unit 0 {
    family inet {
      address 192.168.3.6/24;
    }
    family mpls;
  }
}
ge-0/2/3 {
  unit 0 {
    family inet {
      address 192.168.4.9/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.5/32;
    }
  }
}
```

2. On Router ASBR2, configure the RSVP and MPLS protocols to support the LSP by specifying the Gigabit Ethernet interface that is facing the P2 router.

Configure the OSPF protocol by specifying the Gigabit Ethernet interface that is facing the P2 router and the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```
[edit protocols]
rsvp {
  interface ge-0/2/3.0;
  interface lo0.0;
}
mpls {
  label-switched-path To_PE2 {
    to 192.0.2.7;
  }
  interface lo0.0;
  interface ge-0/2/3.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
```



```

        interface ge-0/2/3.0;
        interface lo0.0;
    }
}

```

3. On Router ASBR2, create the **To-PE2** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE2.

```

[edit protocols]
bgp {
  group To-PE2 {
    type internal;
    local-address 192.0.2.5;
    neighbor 192.0.2.7 {
      family inet-vpn {
        unicast;
      }
    }
  }
}

```

4. On Router ASBR2, create the **To-ASBR1** external BGP peer group. Enable the router to use BGP to advertise NLRI for unicast routes. Specify the neighbor IP peer address as the Gigabit Ethernet interface on Router ASBR1.

```

[edit protocols]
bgp {
  group To-ASBR1 {
    type external;
    family inet-vpn {
      unicast;
    }
    neighbor 192.168.3.8 {
      peer-as 100;
    }
  }
}

```

5. On Router ASBR2, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 200;

```

## Configuring Router P2

### Step-by-Step Procedure

1. On Router P2, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
so-0/0/0 {
  unit 0 {

```

```

        family inet {
            address 192.168.5.10/24;
        }
        family mpls;
    }
}
ge-0/2/2 {
    unit 0 {
        family inet {
            address 192.168.4.11/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.6/32;
        }
    }
}
}

```

2. On Router P2, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
    interface so-0/0/0.0;
    interface ge-0/2/2.0;
    interface lo0.0;
}
mpls {
    interface lo0.0;
    interface ge-0/2/2.0;
    interface so-0/0/0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/2/2.0;
        interface so-0/0/0.0;
        interface lo0.0;
    }
}

```

## Configuring Router PE2

### Step-by-Step Procedure

1. On Router PE2, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET and Fast Ethernet interfaces.

```
[edit interfaces]
so-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.5.12/24;
    }
    family mpls;
  }
}
fe-0/3/1 {
  unit 0 {
    family inet {
      address 192.168.6.13/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.7/32;
    }
  }
}
```

2. On Router PE2, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the BGP peer group within the VRF. Specify AS **20** as the peer AS and specify the IP address of the Fast Ethernet interface on Router CE1 as the neighbor address.

```
[edit routing-instances]
vpn2CE2 {
  instance-type vrf;
  interface fe-0/3/1.0;
  route-distinguisher 1:100;
  vrf-import vpnimport;
  vrf-export vpnexport;
  protocols {
    bgp {
      group To_CE2 {
        peer-as 20;
        neighbor 192.168.6.14;
      }
    }
  }
}
```

3. On Router PE2, configure the RSVP and MPLS protocols to support the LSP. Configure the LSP to ASBR2 and specify the IP address of the logical loopback interface on Router ASBR2. Configure a BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE2. Specify the neighbor address as the logical loopback interface on the Router ASBR2. Specify the **inet-vpn** address family and **unicast** traffic type to enable BGP to carry IPv4 NLRI for VPN routes. Configure the OSPF protocol. Specify the core-facing SONET interface and the logical loopback interface on Router PE2.

```
[edit protocols]
rsvp {
  interface so-0/0/1.0;
  interface lo0.0;
}
mpls {
  label-switched-path To-ASBR2 {
    to 192.0.2.5;
  }
  interface so-0/0/1.0;
  interface lo0.0;
}
bgp {
  group To_ASBR2 {
    type internal;
    local-address 192.0.2.7;
    neighbor 192.0.2.5 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface so-0/0/1.0;
    interface lo0.0;
  }
}
```

4. On Router PE2, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 200;
```

5. On Router PE2, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```
[edit policy-options]
policy-statement vpnexport {
  term 1 {
    from protocol bgp;
    then {
      community add test_comm;
```

```

        accept;
    }
}
term 2 {
    then reject;
}
}

```

6. On Router PE2, configure a policy to import routes from BGP that have the **test\_comm** community attached.

```

[edit policy-options]
policy-statement vpnimport {
    term 1 {
        from {
            protocol bgp;
            community test_comm;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}

```

7. On Router PE1, define the **test\_comm** BGP community with a route target.

```

[edit policy-options]
community test_comm members target:1:100;

```

## Configuring Router CE2

### Step-by-Step Procedure

1. On Router CE2, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE2 and Router PE2. Specify the **inet** address family type.

```

[edit interfaces]
fe-3/0/0 {
    unit 0 {
        family inet {
            address 192.168.6.14/24;
        }
    }
}

```

2. On Router CE2, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.

```

[edit interfaces]
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.8/32;
        }
    }
}

```

```

    }
  }
}

```

- On Router CE2, define a policy named **myroutes** that accepts direct routes.

```

[edit policy-options]
policy-statement myroutes {
  from protocol direct;
  then accept;
}

```

- On Router CE2, configure a routing protocol. The routing protocol can be a static route, RIP, OSPF, ISIS, or EBGp. In this example, we configure EBGp. Specify AS **200** as the peer AS and specify the BGP neighbor IP address as the Fast Ethernet interface of Router PE2. Include the **export** statement.

```

[edit protocols]
bgp {
  group To_PE2 {
    neighbor 192.168.6.13 {
      export myroutes;
      peer-as 200;
    }
  }
}

```

- On Router CE2, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 20;

```

## Verifying the VPN Operation

### Step-by-Step Procedure

- Commit the configuration on each router.



**NOTE:** The MPLS labels shown in this example will be different than the labels used in your configuration.

- On Router PE1, display the routes for the **vpn2CE1** routing instance using the **show ospf route** command. Verify that the **192.0.2.1** route is learned from OSPF.

```
user@PE1> show ospf route instance vpn2CE1
```

Topology default Route Table:

Prefix	Path	Route	NH	Metric	NextHop	NextHop
	Type	Type	Type		Interface	addr/label

```

192.0.2.1          Intra Router      IP      1 fe-1/2/3.0    203.0.113.1
192.0.2.1/32      Intra Network    IP      1 fe-1/2/3.0    203.0.113.1

```

3. On Router PE1, use the **show route advertising-protocol** command to verify that Router PE1 advertises the **192.0.2.1** route to Router ASBR1 using MP-BGP with the VPN MPLS label.

```

user@PE1> show route advertising-protocol bgp 192.0.2.4 extensive
vpn2CE1.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group To_ASBR1 type Internal
    Route Distinguisher: 1:100
    VPN Label: 299952
    Nexthop: Self
    Flags: Nexthop Change
    MED: 1
    Localpref: 100
    AS path: [100] I
    Communities: target:1:100 rte-type:0.0.0.2:1:0

```

4. On Router ASBR1, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.1** route and places it in the **bgp.l3vpn.0** routing table.

```

user@ASBR1> show route receive-protocol bgp 192.0.2.2 extensive
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

inet.3: 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)
* 1:100:192.0.2.1/32 (1 entry, 1 announced)
  Route Distinguisher: 1:100
  VPN Label: 299952
  Nexthop: 192.0.2.2
  MED: 1
  Localpref: 100
  AS path: I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

5. On Router ASBR1, use the **show route advertising-protocol** command to verify that Router ASBR1 advertises the **192.0.2.1** route to Router ASBR2.

```

user@ASBR1> show route advertising-protocol bgp 192.168.3.6 extensive
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:192.0.2.1/32 (1 entry, 1 announced)
  BGP group To-ASBR2 type External
    Route Distinguisher: 1:100
    VPN Label: 299984
    Nexthop: Self
    Flags: Nexthop Change
    AS path: [100] I
    Communities: target:1:100 rte-type:0.0.0.2:1:0

```

6. On Router ASBR2, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.1** route and places it in the **bgp.l3vpn.0** routing table.

```
user@ASBR2> show route receive-protocol bgp 192.168.3.8 extensive
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

inet.3: 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)
* 1:100:192.0.2.1/32 (1 entry, 1 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 299984
  Nexthop: 192.168.3.8
  AS path: 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0
```

7. On Router ASBR2, use the **show route advertising-protocol** command to verify that Router ASBR2 advertises the **192.0.2.1** route to Router PE2.

```
user@ASBR2> show route advertising-protocol bgp 192.0.2.7 extensive
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:192.0.2.1/32 (1 entry, 1 announced)
  BGP group To-PE2 type Internal
  Route Distinguisher: 1:100
  VPN Label: 300048
  Nexthop: Self
  Flags: Nexthop Change
  Localpref: 100
  AS path: [200] 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0
```

8. On Router PE2, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.1** route and places it in the **vpn2CE2.inet.0** routing table.

```
user@PE2> show route receive-protocol bgp 192.0.2.5 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

__juniper_private1__.inet.0: 14 destinations, 14 routes (8 active, 0 holddown,
6 hidden)

__juniper_private2__.inet.0: 1 destinations, 1 routes (0 active, 0 holddown,
1 hidden)

vpn2CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 300048
  Nexthop: 192.0.2.5
  Localpref: 100
  AS path: 100 I
```



```

AS path: Recorded
Communities: target:1:100 rte-type:0.0.0.2:1:0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

* 1:100:192.0.2.1/32 (1 entry, 0 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 300048
  Nexthop: 192.0.2.5
  Localpref: 100
  AS path: 100 I
  AS path: Recorded
  Communities: target:1:100 rte-type:0.0.0.2:1:0

__juniper_private1__.inet6.0: 4 destinations, 4 routes (4 active, 0 holddown,
0 hidden)

```

9. On Router PE2, use the **show route advertising-protocol** command to verify that Router PE2 advertises the **192.0.2.1** route to Router CE2 through the **To\_CE2** peer group.

```

user@PE2> show route advertising-protocol bgp 192.168.6.14 extensive
vpn2CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group To_CE2 type External
  Nexthop: Self
  AS path: [200] 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

10. On Router CE2, use the **show route** command to verify that Router CE2 receives the **192.0.2.1** route from Router PE2.

```

user@CE2> show route 192.0.2.1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32          *[BGP/170] 00:25:36, localpref 100
                     AS path: 200 100 I
                     > to 192.168.6.13 via fe-3/0/0.0

```

11. On Router CE2, use the **ping** command and specify **192.0.2.8** as the source of the ping packets to verify connectivity with Router CE1.

```

user@CE2> ping 192.0.2.1 source 192.0.2.8
PING 192.0.2.1 (192.0.2.1): 56 data bytes
64 bytes from 192.0.2.1: icmp_seq=0 ttl=58 time=4.786 ms
64 bytes from 192.0.2.1: icmp_seq=1 ttl=58 time=10.210 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=58 time=10.588 ms

```

12. On Router PE2, use the **show route** command to verify that the traffic is sent with an inner label of **300048** and a top label of **299776**.

```

user@PE2> show route 192.0.2.1 detail

```

```

vpn2CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
192.0.2.1/32 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 1:100
          Next hop type: Indirect
          Next-hop reference count: 3
          Source: 192.0.2.5
          Next hop type: Router, Next hop index: 653
          Next hop: via so-0/0/1.0 weight 0x1, selected
          Label-switched-path To-PE2
          Label operation: Push 300048, Push 299776(top)
          Protocol next hop: 192.0.2.5
          Push 300048
          Indirect next hop: 8c61138 262143
          State: <Secondary Active Int Ext>
          Local AS: 200 Peer AS: 200
          Age: 27:48      Metric2: 2
          Task: bgp_200.192.0.2.5+60185
          Announcement bits (3): 0-RT 1-KRT 2-BGP RT Background
          AS path: 100 I
          AS path: Recorded
          Communities: target:1:100 rte-type:0.0.0.2:1:0
          Accepted
          VPN Label: 300048
          Localpref: 100
          Router ID: 192.0.2.5
          Primary Routing Table bgp.l3vpn.0

```

13. On Router ASBR2, use the **show route table** command to verify that Router ASBR2 receives the traffic after the top label is popped by Router P2, that label **300048** is swapped with label **299984**, and that the packet is sent toward Router ASBR1 through interface **ge-0/1/1.0**.

```

user@ASBR2> show route table mpls.0 detail
300048 (1 entry, 1 announced)
  *VPN    Preference: 170
          Next hop type: Router, Next hop index: 648
          Next-hop reference count: 2
          Source: 192.168.3.8           Next hop: 192.168.3.8 via
ge-0/1/1.0, selected
          Label operation: Swap 299984
          State: <Active Int Ext>
          Local AS: 200
          Age: 30:39
          Task: BGP RT Background
          Announcement bits (1): 0-KRT
          AS path: 100 I
          Ref Cnt: 1
          Communities: target:1:100 rte-type:0.0.0.2:1:0

```

14. On Router ASBR1, use the **show route table** command to verify that Router ASBR1 receives the traffic with label **299984**, swaps the label with **299952**, and pushes a new top label of **299792**.

```

user@ASBR1> show route table mpls.0 detail
299984 (1 entry, 1 announced)
  *VPN    Preference: 170
          Next hop type: Indirect
          Next-hop reference count: 2

```

```

Source: 192.0.2.2
Next hop type: Router, Next hop index: 538
Next hop: 192.168.2.5 via ge-0/0/0.0 weight 0x1, selected
Label-switched-path To_PE1
Label operation: Swap 299952, Push 299792(top)
Protocol next hop: 192.0.2.2
Swap 299952
Indirect next hop: 8af70a0 262142
State: <Active Int Ext>
Local AS: 100
Age: 34:09 Metric2: 2
Task: BGP RT Background
Announcement bits (1): 0-KRT
AS path: I
Ref Cnt: 1
Communities: target:1:100 rte-type:0.0.0.2:1:0

```

15. On Router PE1, use the **show route table** command to verify that Router PE1 receives the traffic with label **299952**, and then pops the inner label.

```

user@PE1> show route route table mpls.0 detail
299952 (1 entry, 1 announced)
  *VPN Preference: 170
    Next hop type: Router, Next hop index: 536
    Next-hop reference count: 2
    Next hop: 203.0.113.1 via fe-1/2/3.0, selected
    Label operation: Pop
    State: Active Int Ext
    Local AS: 100
    Age: 40:26
    Task: BGP RT Background
    Announcement bits (1): 0-KRT
    AS path: I
    Ref Cnt: 1
    Communities: rte-type:0.0.0.2:1:0

```

**Related Documentation** • [Interprovider VPNs on page 317](#)

## Example: Configuring Interprovider Layer 3 VPN Option C

Interprovider Layer 3 VPN Option C provides interprovider multihop EBGp redistribution of labeled VPN-IPv4 routes between source and destination ASs, with EBGp redistribution of labeled IPv4 routes from AS to neighboring AS. Compared to Option A and Option B, Option C is the most scalable solution. To configure an interprovider Layer 3 VPN option C service, you need to configure the AS border routers and the PE routers connected to the end customer's CE routers using multihop EBGp.

This example provides a step-by-step procedure to configure interprovider layer 3 VPN option C, which is one of the recommended implementations of MPLS VPN when that service is required by a customer that has more than one AS but not all of the customer's

ASs can be serviced by the same service provider (SP). It is organized in the following sections:

- [Requirements on page 376](#)
- [Configuration Overview and Topology on page 376](#)
- [Configuration on page 378](#)

## Requirements

This example requires the following hardware and software components:

- Junos OS Release 9.5 or later.
- Eight Juniper Networks M Series Multiservice Edge Routers, T Series Core Routers, TX Matrix Routers, or MX Series 3D Universal Edge Routers.

## Configuration Overview and Topology

Interprovider layer 3 VPN option C is a very scalable interprovider VPN solution to the problem of providing VPN services to a customer that has different sites, not all of which can use the same SP.

*RFC 4364* section 10, refers to this method as multihop EBGp redistribution of labeled VPN-IPv4 routes between source and destination ASs, with EBGp redistribution of labeled IPv4 routes from AS to neighboring AS.

This solution is similar to the solution described in *Implementing Interprovider Layer 3 VPN Option B*, except internal IPv4 unicast routes are advertised instead of external VPN-IPv4-unicast routes, using EBGp. Internal routes are internal to leaf SPs (SP1 and SP2 in this example), and external routes are those learned from the end customer requesting VPN services.

In this configuration:

- After the loopback address of Router PE2 is learned by Router PE1 and the loopback address of Router PE1 is learned by Router PE2, the end PE routers establish an MP-EBGP session for exchanging VPN-IPv4 routes.
- Since VPN-IPv4 routes are exchanged among end PE routers, any other router on the path from Router PE1 and Router PE2 does not need to keep or install VPN-IPv4 routes in their routing information base (RIB) or forwarding information base (FIB) tables.
- An MPLS path needs to be established between Router PE1 and Router PE2.

*RFC 4364* describes only one solution that uses a BGP labeled-unicast approach. In this approach, the ASBR routers advertise the loopback addresses of the PE routers and associate each prefix with a label according to *RFC 3107*. Service providers may use RSVP or LDP to establish an LSP between ASBR routers and PE routers in their internal network.

In this network, ASBR2 receives label information associated with the loopback IP address of Router PE1 and advertises another label to Router ASBR1 using MP-EBGP labeled-unicast. Meanwhile, the ASBRs build their own MPLS forwarding table according

to the received and advertised routes and labels. Router ASBR1 uses its own IP address as the next-hop information.

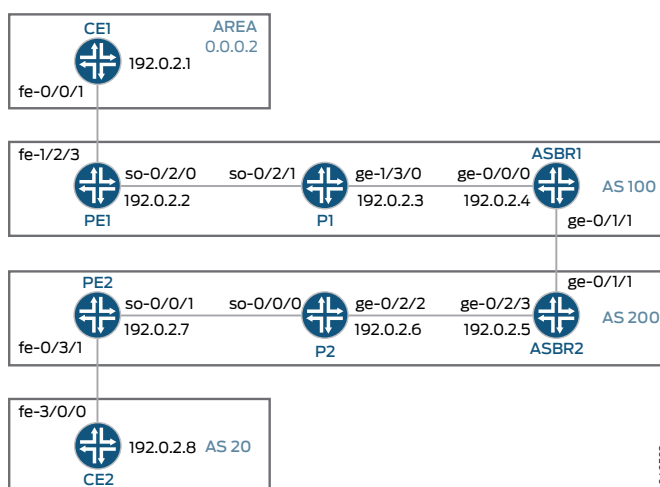
Router ASBR2 receives this prefix associated with a label, assigns another label, changes the next-hop address to its own address, and advertises it to Router PE1. Router PE1 now has an update with the label information and next-hop to Router ASBR1. Also, Router PE1 already has a label associated with the IP address of Router ASBR1. If Router PE1 sends an IP packet to Router PE2, it pushes two labels: one for the IP address of Router PE2 (obtained using MP-IBGP labeled-unicast advertisement) and one for the IP address of Router ASBR1 (obtained using LDP or RSVP).

Router ASBR1 then pops the outer label and swaps the inner label with the label learned from a neighbor ASBR for its neighboring PE router. Router ASBR2 performs a similar function and swaps the incoming label (only one) and pushes another label that is associated with the address of Router PE2. Router PE2 pops both labels and passes the remaining IP packet to its own CPU. After the end-to-end connection among the PE routers is created, the PE routers establish an MP-EBGP session to exchange VPN-IPv4 routes.

In this solution, PE routers push three labels onto the IP packet coming from the VPN end user. The inner-most label, obtained using MP-EBGP, determines the correct VPN routing and forwarding (VRF) routing instance at the remote PE. The middle label is associated with the IP address of the remote PE and is obtained from an ASBR using MP-IBGP labeled-unicast. The outer label is associated with the IP addresses of the ASBRs and is obtained using LDP or RSVP.

The physical topology of the network is shown in [Figure 47 on page 377](#).

**Figure 47: Physical Topology of Interprovider Layer 3 VPN Option C**



## Configuration



**NOTE:** The procedure presented here is written with the assumption that the reader is already familiar with MPLS MVPN configuration. This example focuses on explaining the unique configuration required for carrier-of-carriers solutions for VPN services to different sites.

To configure interprovider layer 3 VPN option C, perform the following tasks:

- [Configuring Router CE1 on page 378](#)
- [Configuring Router PE1 on page 379](#)
- [Configuring Router P1 on page 382](#)
- [Configuring Router ASBR1 on page 383](#)
- [Configuring Router ASBR2 on page 386](#)
- [Configuring Router P2 on page 388](#)
- [Configuring Router PE2 on page 389](#)
- [Configuring Router CE2 on page 392](#)
- [Verifying the VPN Operation on page 393](#)

### Configuring Router CE1

#### Step-by-Step Procedure

1. On Router CE1, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE1 and Router PE1. Specify the **inet** address family type.

```
[edit interfaces fe-0/0/1.0]
family inet {
  address 198.51.100.1/24;
}
```

2. On Router CE1, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.

```
[edit interfaces lo0]
unit 0 {
  family inet {
    address 192.0.2.1/32;
  }
}
```

3. On Router CE1, configure a routing protocol. The routing protocol can be a static route, RIP, OSPF, ISIS, or EBGp. In this example we configure OSPF. Include the logical interface for the link between Router CE1 and Router PE1 and the logical loopback interface of Router CE1.

```
[edit protocols]
ospf {
```

```

    area 0.0.0.2 {
        interface fe-0/0/1.0;
        interface lo0.0 {
            passive;
        }
    }
}

```

### Configuring Router PE1

#### Step-by-Step Procedure

1. On Router PE1, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET interfaces.

```

[edit interfaces]
so-0/2/0 {
    unit 0 {
        family inet {
            address 192.168.1.2/24;
        }
        family mpls;
    }
}
fe-1/2/3 {
    unit 0 {
        family inet {
            address 198.51.100.3/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.2/32;
        }
    }
}

```

2. On Router PE1, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the OSPF protocol within the VRF. Specify the customer-facing Fast Ethernet interface and specify the export policy to export BGP routes into OSPF.

```

[edit routing-instances]
vpn2CE1 {
    instance-type vrf;
    interface fe-1/2/3.0;
    route-distinguisher 1:100;
    vrf-import vpnimport;
    vrf-export vpnexport;
    protocols {

```

```

ospf {
    export bgp-to-ospf;
    area 0.0.0.2 {
        interface fe-1/2/3.0;
    }
}
}
}

```

3. On Router PE1, configure the RSVP and MPLS protocols to support the LSP. Configure the LSP to Router ASBR1 and specify the IP address of the logical loopback interface on Router ASBR1. Configure the OSPF protocol. Specify the core-facing SONET interface and specify the logical loopback interface on Router PE1.

```

[edit protocols]
rsvp {
    interface so-0/2/0.0;
    interface lo0.0;
}
mpls {
    label-switched-path To-ASBR1 {
        to 192.0.2.4;
    }
    interface so-0/2/0.0;
    interface lo0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface so-0/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
}

```

4. On Router PE1, configure the **To\_ASBR1** peer BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE1. Specify the neighbor address as the logical loopback interface on Router ASBR1. Specify the **inet** address family. For a PE router to install a route in the VRF, the next hop must resolve to a route stored within the **inet.3** table. The **labeled-unicast resolve-vpn** statements allow labeled routes to be placed in the **inet.3** routing table for route resolution, which are then resolved for PE router connections where the remote PE is located across another AS.

```

[edit protocols]
bgp {
    group To_ASBR1 {
        type internal;
        local-address 192.0.2.2;
        neighbor 192.0.2.4 {
            family inet {
                labeled-unicast {
                    resolve-vpn;
                }
            }
        }
    }
}

```



```

    }
  }
}

```

5. On Router PE1, configure multihop EBGP toward PE2. Specify the **inet-vpn** family.

```

[edit protocols]
bgp {
  group To_PE2 {
    multihop {
      ttl 20;
    }
    local-address 192.0.2.2;
    family inet-vpn {
      unicast;
    }
    neighbor 192.0.2.7 {
      peer-as 200;
    }
  }
}

```

6. On Router PE1, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 100;

```

7. On Router PE1, configure a policy to export the BGP routes into OSPF.

```

[edit policy-options]
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
  term 2 {
    then reject;
  }
}

```

8. On Router PE1, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```

[edit policy-options]
policy-statement vpnexport {
  term 1 {
    from protocol ospf;
    then {
      community add test_comm;
      accept;
    }
  }
  term 2 {

```

```

        then reject;
    }
}

```

9. On Router PE1, configure a policy to import routes from BGP that have the **test\_comm** community attached.

```

[edit policy-options]
policy-statement vpnimport {
  term 1 {
    from {
      protocol bgp;
      community test_comm;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}

```

10. On Router PE1, define the **test\_comm** BGP community with a route target.

```

[edit policy-options]
community test_comm members target:1:100;

```

### Configuring Router P1

#### Step-by-Step Procedure

1. On Router P1, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
so-0/2/1 {
  unit 0 {
    family inet {
      address 192.168.1.4/24;
    }
    family mpls;
  }
}
ge-1/3/0 {
  unit 0 {
    family inet {
      address 192.168.2.5/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {

```

```

        address 192.0.2.3/32;
    }
}

```

2. On Router P1, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
    interface so-0/2/1.0;
    interface ge-1/3/0.0;
    interface lo0.0;
}
mpls {
    interface lo0.0;
    interface ge-1/3/0.0;
    interface so-0/2/1.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-1/3/0.0;
        interface so-0/2/1.0;
        interface lo0.0 {
            passive;
        }
    }
}

```

### Configuring Router ASBR1

#### Step-by-Step Procedure

1. On Router ASBR1, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** addresses families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
ge-0/0/0 {
    unit 0 {
        family inet {
            address 192.168.2.6/24;
        }
        family mpls;
    }
}
ge-0/1/1 {
    unit 0 {
        family inet {
            address 192.168.3.7/24;
        }
    }
}

```

```

    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.4/32;
    }
  }
}
}

```

2. On Router ASBR1, configure the protocols to support the LSP.

Configure the RSVP protocol by specifying the Gigabit Ethernet interface that is facing the P1 router and the logical loopback interface.

Configure the MPLS protocol by specifying the Gigabit Ethernet interfaces and the logical loopback interface. Include the **traffic-engineering bgp-igp-both-ribs** statement at the **[edit protocols mpls]** hierarchy level.

Configure the OSPF protocol on the Gigabit Ethernet interface facing the P1 router and the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
  interface ge-0/0/0.0;
  interface lo0.0;
}
mpls {
  traffic-engineering bgp-igp-both-ribs;
  label-switched-path To_PE1 {
    to 192.0.2.2;
  }
  interface lo0.0;
  interface ge-0/0/0.0;
  interface ge-0/1/1.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
}

```

3. On Router ASBR1, create the **To-PE1** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the Gigabit Ethernet interface address of Router PE1.

```

[edit protocols]
bgp {

```

```

group To-PE1 {
  type internal;
  local-address 192.0.2.4;
  neighbor 192.0.2.2 {
    family inet {
      labeled-unicast;
    }
    export next-hop-self;
  }
}

```

4. On Router ASBR1, create the **To-ASBR2** external BGP peer group. Enable the router to use BGP to advertise network layer reachability information (NLRI) for unicast routes. Specify the neighbor IP peer address as the Gigabit Ethernet interface address on Router ASBR2.

```

[edit protocols]
group To-ASBR2 {
  type external;
  family inet {
    labeled-unicast;
  }
  export To-ASBR2;
  neighbor 192.168.3.8 {
    peer-as 200;
  }
}

```

5. On Router ASBR1, configure the BGP local autonomous system number.

```

[edit routing-options]
autonomous-system 100;

```

6. On Router ASBR 1, configure a policy to import routes from BGP that match the 192.0.2.2/24 route.

```

[edit policy-options]
policy-statement To-ASBR2 {
  term 1 {
    from {
      route-filter 192.0.2.2/32 exact;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}

```

7. On Router ASBR 1, define a next-hop self policy and apply it to the IBGP sessions.

```

[edit policy-options]
policy-statement next-hop-self {
  then {
    next-hop self;
  }
}

```

```
}
}
```

## Configuring Router ASBR2

### Step-by-Step Procedure

1. On Router ASBR2, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP address for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
ge-0/1/1 {
  unit 0 {
    family inet {
      address 192.168.3.8/24;
    }
    family mpls;
  }
}
ge-0/2/3 {
  unit 0 {
    family inet {
      address 192.168.4.9/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.5/32;
    }
  }
}
```

2. On Router ASBR2, configure the protocols to support the LSP.

Configure the RSVP protocol by specifying the Gigabit Ethernet interface facing the P2 router and the logical loopback interface .

Configure the MPLS protocol by specifying the Gigabit Ethernet interfaces and the logical loopback interface. Include the **traffic-engineering bgp-igp-both-ribs** statement at the **[edit protocols mpls]** hierarchy level.

Configure the OSPF protocol on the Gigabit Ethernet interface facing the P2 router and the logical loopback interface . Enable OSPF to support traffic engineering extensions.

```
[edit protocols]
rsvp {
  interface ge-0/2/3.0;
  interface lo0.0;
}
mpls {
```

```

traffic-engineering bgp-igp-both-ribs;
label-switched-path To_PE2 {
  to 192.0.2.7;
}
interface lo0.0
interface ge-0/2/3.0;
interface ge-0/1/1.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/2/3.0;
    interface lo0.0 {
      passive;
    }
  }
}
}

```

3. On Router ASBR2, create the **To-PE2** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE2.

```

[edit protocols]
bgp {
  group To-PE2 {
    type internal;
    local-address 192.0.2.5;
    export next-hop-self;
    neighbor 192.0.2.7 {
      family inet {
        labeled-unicast;
      }
      export next-hop-self;
    }
  }
}

```

4. On Router ASBR2, create the **To-ASBR1** external BGP peer group. Enable the router to use BGP to advertise NLRI for unicast routes. Specify the neighbor IP peer address as the Gigabit Ethernet interface address on Router ASBR1.

```

[edit protocols]
bgp {
  group To-ASBR1 {
    type external;
    family inet {
      labeled-unicast;
    }
    export To-ASBR1;
    neighbor 192.168.3.7 {
      peer-as 100;
    }
  }
}

```

5. On Router ASBR2 configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 200;
```

6. On Router ASBR2, configure a policy to import routes from BGP that match the 192.0.2.7/24 route.

```
[edit policy-options]
policy-statement To-ASBR1 {
  term 1 {
    from {
      route-filter 192.0.2.7/32 exact;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

7. On Router ASBR 2, define a next-hop self policy.

```
[edit policy-options]
policy-statement next-hop-self {
  then {
    next-hop self;
  }
}
```

## Configuring Router P2

### Step-by-Step Procedure

1. On Router P2, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** addresses families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.5.10/24;
    }
    family mpls;
  }
}
ge-0/2/2 {
  unit 0 {
    family inet {
      address 192.168.4.11/24;
    }
    family mpls;
  }
}
```



```

lo0 {
  unit 0 {
    family inet {
      address 192.0.2.6/32;
    }
  }
}

```

2. On Router P2, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
  interface so-0/0/0.0;
  interface ge-0/2/2.0;
  interface lo0.0;
}
mpls {
  interface lo0.0;
  interface ge-0/2/2.0;
  interface so-0/0/0.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/2/2.0;
    interface so-0/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

```

### Configuring Router PE2

#### Step-by-Step Procedure

1. On Router PE2, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET interface.

```

[edit interfaces]
so-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.5.12/24;
    }
    family mpls;
  }
}
fe-0/3/1 {
  unit 0 {

```

```

        family inet {
            address 192.168.6.13/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.7/32;
        }
    }
}
}

```

2. On Router PE2, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the BGP peer group within the VRF. Specify AS **20** as the peer AS and specify the IP address of the Fast Ethernet interface on Router CE1 as the neighbor address.

```

[edit routing-instances]
vpn2CE2 {
    instance-type vrf;
    interface fe-0/3/1.0;
    route-distinguisher 1:100;
    vrf-import vpnimport;
    vrf-export vpnexport;
    protocols {
        bgp {
            group To_CE2 {
                peer-as 20;
                neighbor 192.168.6.14;
            }
        }
    }
}

```

3. On Router PE2, configure the RSVP and MPLS protocols to support the LSP. Configure the LSP to ASBR2 and specify the IP address of the logical loopback interface on Router ASBR2. Configure the OSPF protocol. Specify the core-facing SONET interface and specify the logical loopback interface on Router PE2.

```

[edit protocols]
rsvp {
    interface so-0/0/1.0;
    interface lo0.0;
}
mpls {
    label-switched-path To-ASBR2 {
        to 192.0.2.5;
    }
    interface so-0/0/1.0;
    interface lo0.0;
}

```

```
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface so-0/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
```

4. On Router PE2, configure the **To\_ASBR2** BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE2. Specify the neighbor address as the logical loopback interface on the Router ASBR2.

```
[edit protocols]
bgp {
  group To_ASBR2 {
    type internal;
    local-address 192.0.2.7;
    neighbor 192.0.2.5 {
      family inet {
        labeled-unicast {
          resolve-vpn;
        }
      }
    }
  }
}
```

5. On Router PE2, configure multihop EBGp towards Router PE1 Specify the **inet-vpn** address family.

```
[edit protocols]
bgp {
  group To_PE1 {
    type external;
    local-address 192.0.2.7;
    multihop {
      ttl 20;
    }
    family inet-vpn {
      unicast;
    }
    neighbor 192.0.2.2 {
      peer-as 100;
    }
  }
}
```

6. On Router PE2, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 200;
```

7. On Router PE2, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```
[edit policy-options]
policy-statement vpnexport {
  term 1 {
    from protocol bgp;
    then {
      community add test_comm;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}
```

8. On Router PE2, configure a policy to import routes from BGP that have the **test\_comm** community attached.

```
[edit policy-options]
policy-statement vpnimport {
  term 1 {
    from {
      protocol bgp;
      community test_comm;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

9. On Router PE1, define the **test\_comm** BGP community with a route target.

```
[edit policy-options]
community test_comm members target:1:100;
```

---

### Configuring Router CE2

---

#### Step-by-Step Procedure

1. On Router CE2, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE2 and Router PE2. Specify the **inet** address family type.

```
[edit interfaces]
fe-3/0/0 {
  unit 0 {
    family inet {
      address 192.168.6.14/24;
    }
  }
}
```

2. On Router CE2, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.

```
[edit interfaces lo0]
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.8/32;
    }
  }
}
```

3. On Router CE2, define a policy named **myroutes** that accepts direct routes.

```
[edit policy-options]
policy-statement myroutes {
  from protocol direct;
  then accept;
}
```

4. On Router CE2, configure a routing protocol. The routing protocol can be a static route, RIP, OSPF, ISIS, or EBGp. In this example, we configure EBGp. Specify the BGP neighbor IP address as the logical loopback interface of Router PE1. Apply the **myroutes** policy.

```
[edit protocols]
bgp {
  group To_PE2 {
    neighbor 198.51.100.13 {
      export myroutes;
      peer-as 200;
    }
  }
}
```

5. On Router CE2, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 20;
```

## Verifying the VPN Operation

### Step-by-Step Procedure

1. Commit the configuration on each router.



**NOTE:** The MPLS labels shown in this example will be different than the labels used in your configuration.

2. On Router PE1, display the routes for the **vpn2CE1** routing instance using the **show ospf route** command. Verify that the **192.0.2.1** route is learned from OSPF.

```
user@PE1> show ospf route instance vpn2CE1
Topology default Route Table:
```

Prefix	Path	Route	NH	Metric	NextHop	Nexthop
	Type	Type	Type		Interface	addr/label
192.0.2.1	Intra	Router	IP		1 fe-1/2/3.0	
198.51.100.1						
192.0.2.1/32	Intra	Network	IP	1	fe-1/2/3.0	198.51.100.1
198.51.100.0/24	Intra	Network	IP		1 fe-1/2/3.0	

- On Router PE1, use the **show route advertising-protocol** command to verify that Router PE1 advertises the **192.0.2.1** route to Router PE2 using MP-BGP with the VPN MPLS label.

```
user@PE1> show route advertising-protocol bgp 192.0.2.7 extensive
bgp.13vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:192.0.2.1/32 (1 entry, 1 announced)
  BGP group To_PE2 type External
    Route Distinguisher: 1:100
    VPN Label: 300016
    Nexthop: Self
    Flags: Nexthop Change
    MED: 1
    AS path: [100] I
    Communities: target:1:100 rte-type:0.0.0.2:1:0
```

- On Router ASBR1, use the **show route advertising-protocol** command to verify that Router ASBR1 advertises the **192.0.2.2** route to Router ASBR2.

```
user@ASBR1> show route advertising-protocol bgp 192.168.3.8 extensive
inet.0: 14 destinations, 16 routes (14 active, 0 holddown, 0 hidden)
* 192.0.2.2/32 (2 entries, 1 announced)
  BGP group To-PE2 type External
    Route Label: 300172
    Nexthop: Self
    Flags: Nexthop Change
    MED: 2
    AS path: [100] I
```

- On Router ASBR2, use the **show route receive-protocol** command to verify that the router receives and accepts the **192.0.2.2** route.

```
user@ASBR2> show route receive-protocol bgp 192.168.3.7 extensive
inet.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)
* 192.0.2.2/32 (1 entry, 1 announced)
  Accepted
    Route Label: 300172
    Nexthop: 192.168.3.7
    MED: 2
    AS path: 100 I
```

- On Router ASBR2, use the **show route advertising-protocol** command to verify that Router ASBR2 advertises the **192.0.2.2** route to Router PE2.

```
user@ASBR2> show route advertising-protocol bgp 192.0.2.7 extensive
inet.0: 10 destinations, 11 routes (10 active, 0 holddown, 0 hidden)
* 192.0.2.2/32 (1 entry, 1 announced)
```

```

BGP group To-PE2 type Internal
Route Label: 300192
Nexthop: Self
Flags: Nexthop Change
MED: 2
Localpref: 100
AS path: [200] 100 I

```

7. On Router PE2, use the **show route receive-protocol** command to verify that Router PE2 receives the route and puts it in the **inet.0** routing table. Verify that Router PE2 also receives the update from Router PE1 and accepts the route.

```

user@PE2> show route receive-protocol bgp 192.0.2.5 extensive
inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
* 192.0.2.2/32 (1 entry, 1 announced)
  Accepted
  Route Label: 300192
  Nexthop: 192.0.2.5
  MED: 2
  Localpref: 100
  AS path: 100 I
  AS path: Recorded

```

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```

* 192.0.2.2/32 (1 entry, 1 announced)
  Accepted
  Route Label: 300192
  Nexthop: 192.0.2.5
  MED: 2
  Localpref: 100
  AS path: 100 I
  AS path: Recorded

```

8. On Router PE2, use the **show route receive-protocol** command to verify that Router PE2 puts the route in the routing table of the **vpn2CE2** routing instance and advertises the route to Router CE2 using EBGp.

```

user@PE2> show route receive-protocol bgp 192.0.2.2 detail
inet.0: 17 destinations, 18 routes (17 active, 0 holddown, 0 hidden)

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

__juniper_private1__.inet.0: 14 destinations, 14 routes (8 active, 0 holddown,
6 hidden)

__juniper_private2__.inet.0: 1 destinations, 1 routes (0 active, 0 holddown,
1 hidden)

vpn2CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 300016
  Nexthop: 192.0.2.2
  MED: 1
  AS path: 100 I
  AS path: Recorded
  Communities: target:1:100 rte-type:0.0.0.2:1:0

```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 1:100:192.0.2.1/32 (1 entry, 0 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 300016
  Nexthop: 192.0.2.2
  MED: 1
  AS path: 100 I
  AS path: Recorded
  Communities: target:1:100 rte-type:0.0.0.2:1:0

__juniper_private1__.inet6.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

9. On Router PE2, use the **show route advertising-protocol** command to verify that Router PE2 advertises the **192.0.2.1** route to Router CE2 through the **vpn2CE2** peer group.

```
user@PE2> show route advertising-protocol bgp 192.168.6.14 extensive
vpn2CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
* 192.0.2.1/32 (1 entry, 1 announced)
  BGP group vpn2CE2 type External
  Nexthop: Self
  AS path: [200] 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0
```

10. On Router CE2, use the **show route** command to verify that Router CE2 receives the **192.0.2.1** route from Router PE2.

```
user@CE2> show route 192.0.2.1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.1/32          *[BGP/170] 00:25:36, localpref 100
                     AS path: 200 100 I
                     > to 192.168.6.13 via fe-3/0/0.0
```

11. On Router CE2, use the **ping** command and specify **192.0.2.8** as the source of the ping packets to verify connectivity with Router CE1.

```
user@CE2> ping 192.0.2.1 source 192.0.2.8
PING 192.0.2.1 (192.0.2.1): 56 data bytes
64 bytes from 192.0.2.1: icmp_seq=0 ttl=58 time=4.786 ms
64 bytes from 192.0.2.1: icmp_seq=1 ttl=58 time=10.210 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=58 time=10.588 ms
```

12. On Router PE2, use the **show route** command to verify that the traffic is sent with an inner label of **300016**, a middle label of **300192**, and a top label of **299776**.

```
user@PE2> show route 192.0.2.1 detail
vpn2CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
192.0.2.1/32 (1 entry, 1 announced)
```



```

*BGP    Preference: 170/-101
        Route Distinguisher: 1:100
        Next hop type: Indirect
        Next-hop reference count: 3
        Source: 192.0.2.2
        Next hop type: Router, Next hop index: 653
        Next hop: via so-0/0/1.0 weight 0x1, selected
        Label-switched-path To-ASBR2
Label operation: Push 300016, Push 300192, Push 299776(top)
Protocol next hop: 192.0.2.2
Push 300016
Indirect next hop: 8c61138 262142
State: <Secondary Active Ext>
Local AS: 200 Peer AS: 100
Age: 17:33 Metric: 1 Metric2: 2
Task: BGP_100.192.0.2.2+62319
Announcement bits (3): 0-RT 1-KRT 2-BGP RT Background
AS path: 100 I
AS path: Recorded
Communities: target:1:100 rte-type:0.0.0.2:1:0
Accepted
VPN Label: 300016
Localpref: 100
Router ID: 192.0.2.2
Primary Routing Table bgp.l3vpn.0

```

13. On Router ASBR2, use the **show route table** command to verify that Router ASBR2 receives the traffic after the top label is popped by Router P2. Verify that label **300192** is swapped with label **300176** and the traffic is sent towards Router ASBR1 using interface **ge-0/1/1.0**. At this point, the bottom label **300016** is preserved.

```

user@ASBR2# show route table mpls.0 detail
300192 (1 entry, 1 announced)
  *VPN    Preference: 170
          Next hop type: Router, Next hop index: 660
          Next-hop reference count: 2
          Source: 192.168.3.7           Next hop: 192.168.3.7 via
ge-0/1/1.0, selected
          Label operation: Swap 300176
          State: <Active Int Ext>
          Local AS: 200
          Age: 24:01
          Task: BGP RT Background
          Announcement bits (1): 0-KRT
          AS path: 100 I
          Ref Cnt: 1

```

14. On Router ASBR1, use the **show route table** command to verify that when Router ASBR1 receives traffic with label **300176**, it swaps the label with **299824** to reach Router PE1.

```

user@ASBR1> show route table mpls.0 detail
300176 (1 entry, 1 announced)
  *VPN    Preference: 170
          Next hop type: Router, Next hop index: 651
          Next-hop reference count: 2
          Next hop: 192.168.2.5 via ge-0/0/0.0 weight 0x1, selected
          Label operation: Swap 299824
          State: <Active Int Ext>

```

```
Local AS: 100
Age: 25:53
Task: BGP RT Background
Announcement bits (1): 0-KRT
AS path: I
Ref Cnt: 1
```

15. On Router PE1, use the **show route table** command to verify that Router PE1 receives the traffic after the top label is popped by Router P1. Verify that label **300016** is popped and the traffic is sent towards Router CE1 using interface **fe-1/2/3.0**.

```
user@PE1> show route table mpls.0 detail
300016 (1 entry, 1 announced)
  *VPN      Preference: 170
            Next hop type: Router, Next hop index: 643
            Next-hop reference count: 2
            Next hop: 198.51.100.1 via fe-1/2/3.0, selected
            Label operation: Pop
            State:< Active Int Ext>
            Local AS: 100
            Age: 27:37
            Task: BGP RT Background
            Announcement bits (1): 0-KRT
            AS path: I
            Ref Cnt: 1
            Communities: rte-type:0.0.0.2:1:0
```

**Related Documentation** • [Interprovider VPNs on page 317](#)

## CHAPTER 23

# Configuring Carrier-of-Carrier VPNs

- [Understanding Carrier-of-Carriers VPNs on page 400](#)
- [Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service on page 402](#)
- [Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service on page 407](#)
- [Carrier-of-Carriers VPN Example—Customer Provides Internet Service on page 415](#)
- [Carrier-of-Carriers VPN Example—Customer Provides VPN Service on page 425](#)
- [Multiple Instances for LDP and Carrier-of-Carriers VPNs on page 436](#)

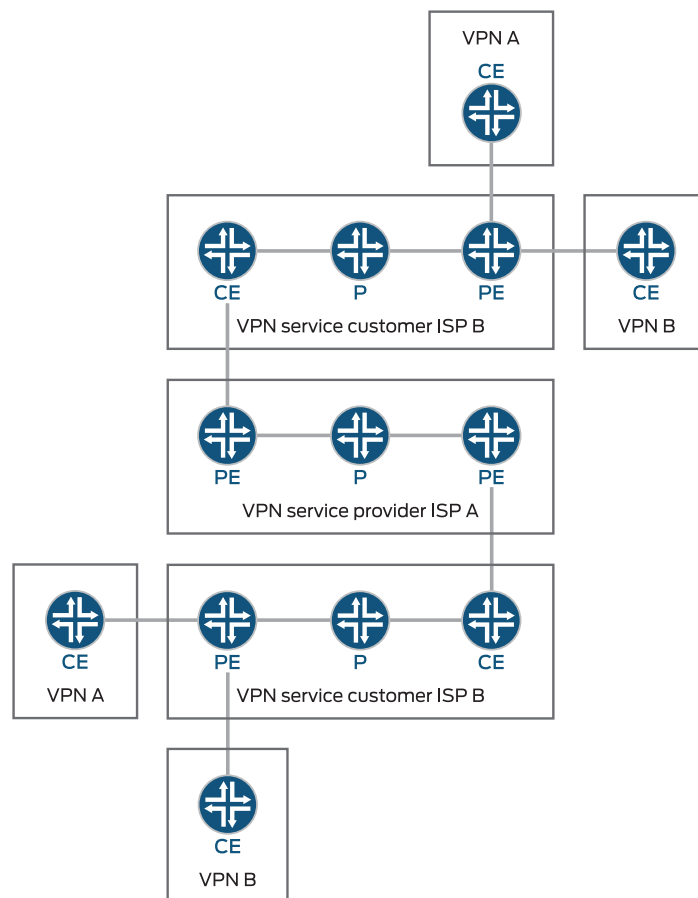
## Understanding Carrier-of-Carriers VPNs

The customer of a VPN service provider might be a service provider for the end customer. The following are the two main types of carrier-of-carriers VPNs (as described in RFC 4364):

- “Internet Service Provider as the Customer” on page 401—The VPN customer is an ISP that uses the VPN service provider’s network to connect its geographically disparate regional networks. The customer does not have to configure MPLS within its regional networks.
- “VPN Service Provider as the Customer” on page 401—The VPN customer is itself a VPN service provider offering VPN service to its customers. The carrier-of-carriers VPN service customer relies on the backbone VPN service provider for inter-site connectivity. The customer VPN service provider is required to run MPLS within its regional networks.

Figure 48 on page 400 illustrates the network architecture used for a carrier-of-carriers VPN service.

Figure 48: Carrier-of-Carriers VPN Architecture



807197

This topic covers the following:

- [Internet Service Provider as the Customer on page 401](#)
- [VPN Service Provider as the Customer on page 401](#)

## Internet Service Provider as the Customer

In this type of carrier-of-carriers VPN configuration, ISP A configures its network to provide Internet service to ISP B. ISP B provides the connection to the customer wanting Internet service, but the actual Internet service is provided by ISP A.

This type of carrier-of-carriers VPN configuration has the following characteristics:

- The carrier-of-carriers VPN service customer (ISP B) does not need to configure MPLS on its network.
- The carrier-of-carriers VPN service provider (ISP A) must configure MPLS on its network.
- MPLS must also be configured on the CE routers and PE routers connected together in the carrier-of-carriers VPN service customer's and carrier-of-carriers VPN service provider's networks.

## VPN Service Provider as the Customer

A VPN service provider can have customers that are themselves VPN service providers. In this type of configuration, also called a hierarchical or recursive VPN, the customer VPN service provider's VPN-IPv4 routes are considered external routes, and the backbone VPN service provider does not import them into its VRF table. The backbone VPN service provider imports only the customer VPN service provider's internal routes into its VRF table.

The similarities and differences between interprovider and carrier-of-carriers VPNs are shown in [Table 10 on page 401](#).

**Table 10: Comparison of Interprovider and Carrier-of-Carriers VPNs**

Feature	ISP Customer	VPN Service Provider Customer
Customer edge device	AS border router	PE router
IBGP sessions	Carry IPv4 routes	Carry external VPN-IPv4 routes with associated labels
Forwarding within the customer network	MPLS is optional	MPLS is required

Support for VPN service as the customer is supported on QFX10000 switches starting with Junos OS Release 17.1R1.

## Release History Table

Release	Description
17.1R1	Support for VPN service as the customer is supported on QFX10000 switches starting with Junos OS Release 17.1R1.

## Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)
- [Interprovider VPNs on page 317](#)

## Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service

You can configure a carrier-of-carriers VPN service for customers who want to provide basic Internet service. The carrier-of-carriers VPN service provider must configure MPLS in its network, although this configuration is optional for the carrier service customer. [Figure 48 on page 400](#) shows how the routers or switches in this type of service interconnect.

To configure a carrier-of-carriers VPN, perform the tasks described in the following sections:

- [Configuring the Carrier-of-Carriers VPN Service Customer's CE Router on page 402](#)
- [Configuring the Carrier-of-Carriers VPN Service Provider's PE Routers on page 404](#)

### Configuring the Carrier-of-Carriers VPN Service Customer's CE Router

The carrier-of-carriers VPN service customer's router (or switch) acts as a CE router with respect to the service provider's PE router or switch. The following sections describe how to configure the carrier-of-carriers VPN service customer's CE router or switch:

- [Configuring MPLS on page 402](#)
- [Configuring BGP on page 403](#)
- [Configuring OSPF on page 403](#)
- [Configuring Policy Options on page 404](#)

#### Configuring MPLS

To configure MPLS on the customer's CE router or switch, include the **mpls** statement:

```
mpls {
  traffic-engineering bgp-igp;
  interface interface-name;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

## Configuring BGP

To configure a group to collate the customer's internal routes, include the **bgp** statement:

```
bgp {
  group group-name {
    type internal;
    local-address address;
    neighbor address;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

The customer's CE router (or switch) must be able to send labels to the VPN service provider's router. Enable this by including the **labeled-unicast** statement in the configuration for the BGP group:

```
bgp {
  group group-name {
    export internal;
    peer-as as-number;
    neighbor address {
      family inet {
        labeled-unicast;
      }
    }
  }
}
```

You can include the **bgp** statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

## Configuring OSPF

To configure OSPF on the customer's CE router or switch, include the **ospf** statement:

```
ospf {
  area area-id {
    interface interface-name {
      passive;
    }
    interface interface-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]

- [\[edit logical-systems \*logical-system-name\* protocols\]](#)

### Configuring Policy Options

---

To configure policy options on the customer's CE router or switch, include the **policy-statement** statement:

```
policy-statement statement-name {  
  term term-name {  
    from protocol [ospf direct ldp];  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [\[edit policy-options\]](#)
- [\[edit logical-systems \*logical-system-name\* policy-options\]](#)

## Configuring the Carrier-of-Carriers VPN Service Provider's PE Routers

The service provider's PE routers connect to the customer's CE routers and forward the customer's VPN traffic across the provider's network.

The following sections describe how to configure the carrier-of-carriers VPN service provider's PE routers:

- [Configuring MPLS on page 404](#)
- [Configuring BGP on page 405](#)
- [Configuring IS-IS on page 405](#)
- [Configuring LDP on page 405](#)
- [Configuring a Routing Instance on page 406](#)
- [Configuring Policy Options on page 406](#)

### Configuring MPLS

---

To configure MPLS on the provider's PE routers or switches include the **mpls** statement:

```
mpls {  
  interface interface-name;  
  interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- [\[edit protocols\]](#)
- [\[edit logical-systems \*logical-system-name\* protocols\]](#)



---

## Configuring BGP

To configure a BGP session with the provider PE router at the other end of the provider's network, include the **bgp** statement:

```
bgp {  
  group group-name {  
    type internal;  
    local-address address;  
    family inet-vpn {  
      any;  
    }  
    neighbor address;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

---

## Configuring IS-IS

To configure IS-IS on the provider's PE routers or switches, include the **isis** statement:

```
isis {  
  interface interface-name;  
  interface interface-name {  
    passive;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

---

## Configuring LDP

To configure LDP on the provider's PE routers or switches, include the **ldp** statement:

```
ldp {  
  interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

## Configuring a Routing Instance

To configure Layer 3 VPN service with the customer's CE router or switch, include the **labeled-unicast** statement in the configuration for the routing instance so the PE router (or switch) can send labels to the customer's CE router or switch:

```

routing-instance-name {
  instance-type vrf;
  interface interface-name;
  route-distinguisher address;
  vrf-import policy-name;
  vrf-export policy-name;
  protocols {
    bgp {
      group group-name {
        peer-as as-number;
        neighbor address {
          family inet {
            labeled-unicast;
          }
        }
      }
    }
  }
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances]
- [edit logical-systems *logical-system-name* routing-instances]

## Configuring Policy Options

To configure a policy statement to import routes from the customer's CE router or switch, include the **policy-statement** statement:

```

policy-statement policy-name {
  term term-name {
    from {
      protocol bgp;
      community community-name;
    }
    then accept;
  }
  term term-name {
    then reject;
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

To configure a policy statement to export routes to the customer's CE router or switch, include the **policy-statement** and **community** statements:

```
policy-statement policy-name {
  term term-name {
    from protocol bgp;
    then {
      community add community-name;
      accept;
    }
  }
  term term-name {
    then reject;
  }
}
community community-name members value;
```

You can include these statements at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

#### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)

## Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service

You can configure a carrier-of-carriers VPN service for customers who want VPN service.

To configure the routers (or switches) in the customer's and provider's networks to enable carrier-of-carriers VPN service, perform the steps in the following sections:

- [Configuring the Carrier-of-Carriers Customer's PE Router on page 407](#)
- [Configuring the Carrier-of-Carriers Customer's CE Router \(or switch\) on page 410](#)
- [Configuring the Provider's PE Router or Switch on page 412](#)

### Configuring the Carrier-of-Carriers Customer's PE Router

The carrier-of-carriers customer's PE router (or switch) is connected to the end customer's CE router (or switch).

The following sections describe how to configure the carrier-of-carriers customer's PE router (or switch):

- [Configuring MPLS on page 408](#)
- [Configuring BGP on page 408](#)
- [Configuring OSPF on page 408](#)
- [Configuring LDP on page 409](#)

- [Configuring VPN Service in the Routing Instance on page 409](#)
- [Configuring Policy Options on page 410](#)

---

### Configuring MPLS

To configure MPLS on the carrier-of-carriers customer's PE router (or switch), include the **mpls** statement:

```
mpls {  
    interface interface-name;  
    interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

---

### Configuring BGP

Include the **labeled-unicast** statement in the configuration for the IBGP session to the carrier-of-carriers customer's CE router (or switch) (see "[Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service](#)" on page 407), and include the **family-inet-vpn** statement in the configuration for the IBGP session to the carrier-of-carriers PE router (or switch) on the other side of the network:

```
bgp {  
    group group-name {  
        type internal;  
        local-address address;  
        neighbor address {  
            family inet {  
                labeled-unicast;  
                resolve-vpn;  
            }  
        }  
    }  
    neighbor address {  
        family inet-vpn {  
            any;  
        }  
    }  
}
```

You can include these statements at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

---

### Configuring OSPF

To configure OSPF on the carrier-of-carriers customer's PE router (or switch), include the **ospf** statement:

```
ospf {
```

```

    area area-id {
        interface interface-name {
            passive;
        }
        interface interface-name;
    }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

### Configuring LDP

To configure LDP on the carrier-of-carriers customer's PE router (or switch), include the **ldp** statement:

```

ldp {
    interface interface-name;
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

### Configuring VPN Service in the Routing Instance

To configure VPN service for the end customer's CE router (or switch) on the carrier-of-carriers customer's PE router (or switch), include the following statements:

```

instance-type vrf;
interface interface-name;
route-distinguisher address;
vrf-import policy-name;
vrf-export policy-name;
protocols {
    bgp {
        group group-name {
            peer-as as-number;
            neighbor address;
        }
    }
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

## Configuring Policy Options

---

To configure policy options to import and export routes to and from the end customer's CE router (or switch), include the **policy-statement** and **community** statements:

```
policy-statement policy-name {  
  term term-name {  
    from {  
      protocol bgp;  
      community community-name;  
    }  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}  
policy-statement policy-name {  
  term term-name {  
    from protocol bgp;  
    then {  
      community add community-name;  
      accept;  
    }  
  }  
  term term-name {  
    then reject;  
  }  
}  
community community-name members value;
```

You can include these statements at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems *logical-system-name* policy-options]**

## Configuring the Carrier-of-Carriers Customer's CE Router (or switch)

The carrier-of-carriers customer's CE router (or switch) connects to the provider's PE router (or switch). Complete the instructions in the following sections to configure the carrier-of-carriers customers' CE router (or switch):

- [Configuring MPLS on page 410](#)
- [Configuring BGP on page 411](#)
- [Configuring OSPF and LDP on page 412](#)
- [Configuring Policy Options on page 412](#)

## Configuring MPLS

---

In the MPLS configuration for the carrier-of-carriers customer's CE router (or switch), include the interfaces to the provider's PE router (or switch) and to a P router (or switch) in the customer's network:

```

mpls {
  traffic-engineering bgp-igp;
  interface interface-name;
  interface interface-name;
}

```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

### Configuring BGP

In the BGP configuration for the carrier-of-carriers customer's CE router (or switch), configure a group that includes the **labeled-unicast** statement to extend VPN service to the PE router (or switch) connected to the end customer's CE router (or switch):

```

bgp {
  group group-name {
    type internal;
    local-address address;
    neighbor address {
      family inet {
        labeled-unicast;
      }
    }
  }
}

```

You can include the **bgp** statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

To configure a group to send labeled internal routes to the provider's PE router (or switch), include the **bgp** statement:

```

bgp {
  group group-name {
    export internal;
    peer-as as-number;
    neighbor address {
      family inet {
        labeled-unicast;
      }
    }
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

### Configuring OSPF and LDP

---

To configure OSPF and LDP on the carrier-of-carriers customer's CE router (or switch), include the **ospf** and **ldp** statements:

```
ospf {  
  area area-id {  
    interface interface-name {  
      passive;  
    }  
    interface interface-name;  
  }  
}  
ldp {  
  interface interface-name;  
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

### Configuring Policy Options

---

To configure the policy options on the carrier-of-carriers customer's CE router (or switch), include the **policy-statement** statement:

```
policy-statement policy-statement-name {  
  term term-name {  
    from protocol [ ospf direct ldp ];  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

## Configuring the Provider's PE Router or Switch

The carrier-of-carriers provider's PE routers (or switches) connect to the carrier customer's CE routers (or switches) . Complete the instructions in the following sections to configure the provider's PE router (or switch):

- [Configuring MPLS on page 413](#)
- [Configuring a PE-to-PE BGP Session on page 413](#)
- [Configuring IS-IS and LDP on page 413](#)



- [Configuring Policy Options on page 414](#)
- [Configuring a Routing Instance to Send Routes to the CE Router on page 414](#)

### Configuring MPLS

In the MPLS configuration, specify at least two interfaces—one to the customer's CE router (or switch) and one to connect to the provider's PE router (or switch) on the other side of the provider's network:

```
interface interface-name;
interface interface-name;
```

You can include these statements at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

### Configuring a PE-to-PE BGP Session

To configure a PE-to-PE BGP session on the provider's PE routers (or switches) to allow VPN-IPv4 routes to pass between the PE routers (or switches), include the **bgp** statement:

```
bgp {
  group group-name {
    type internal;
    local-address address;
    family inet-vpn {
      any;
    }
    neighbor address;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

### Configuring IS-IS and LDP

To configure IS-IS and LDP on the provider's PE routers (or switches), include the **isis** and **ldp** statements:

```
isis {
  interface interface-name;
  interface interface-name {
    passive;
  }
}
ldp {
  interface interface-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

---

### Configuring Policy Options

To configure policy statements on the provider's PE router (or switch) to export routes to and import routes from the carrier customer's network, include the **policy-statement** and **community** statements:

```
policy-statement statement-name {
  term term-name {
    from {
      protocol bgp;
      community community-name;
    }
    then accept;
  }
  term term-name {
    then reject;
  }
}
policy-statement statement-name {
  term term-name {
    from protocol bgp;
    then {
      community add community-name;
      accept;
    }
  }
  term term-name {
    then reject;
  }
}
community community-name members value;
```

You can include these statements at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

---

### Configuring a Routing Instance to Send Routes to the CE Router

To configure the routing instance on the provider's PE router (or switch) to send labeled routes to the carrier customer's CE router (or switch), include the following statements:

```
instance-type vrf;
interface interface-name;
route-distinguisher value;
vrf-import policy-name;
vrf-export policy-name;
protocols {
  bgp {
```

```

group group-name {
  peer-as as-number;
  neighbor address {
    family inet {
      labeled-unicast;
    }
  }
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

#### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)

## Carrier-of-Carriers VPN Example—Customer Provides Internet Service

In this example, the carrier customer is not required to configure MPLS and LDP on its network. However, the carrier provider must configure MPLS and LDP on its network.

For configuration information see the following sections:


- [Network Topology for Carrier-of-Carriers Service on page 415](#)
- [Configuration for Router A on page 416](#)
- [Configuration for Router B on page 416](#)
- [Configuration for Router C on page 417](#)
- [Configuration for Router D on page 418](#)
- [Configuration for Router E on page 419](#)
- [Configuration for Router F on page 420](#)
- [Configuration for Router G on page 420](#)
- [Configuration for Router H on page 421](#)
- [Configuration for Router I on page 422](#)
- [Configuration for Router J on page 423](#)
- [Configuration for Router K on page 423](#)
- [Configuration for Router L on page 424](#)

### Network Topology for Carrier-of-Carriers Service

A carrier-of-carriers service allows an Internet service provider (ISP) to connect to a transparent outsourced backbone at multiple locations.

[Figure 49 on page 416](#) shows the network topology in this carrier-of-carriers example.

Figure 49: Carrier-of-Carriers VPN Example Network Topology



### Configuration for Router A

In this example, Router A represents an end customer. You configure this router as a CE device.

```
[edit]
protocols {
  bgp {
    group to-routerB {
      export attached;
      peer-as 21;
      as-override;
      neighbor 192.168.197.169;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
```

### Configuration for Router B

Router B can act as the gateway router, responsible for aggregating end customers and connecting them to the network. If a full-mesh IBGP session is configured, you can use route reflectors.

```
[edit]
protocols {
  bgp {
    group int {
      type internal;
      local-address 10.255.14.179;
    }
  }
}
```

```

neighbor 10.255.14.175;
neighbor 10.255.14.181;
neighbor 10.255.14.176;
neighbor 10.255.14.178;
neighbor 10.255.14.177;
}
group to-vpn-blue {
peer-as 1;
neighbor 192.168.197.170;
}
}
ospf {
area 0.0.0.0 {
interface lo0.0 {
passive;
}
interface fe-1/0/3.0;
interface fe-1/0/2.0 {
passive;
}
}
}
}
}

```

## Configuration for Router C

Configure Router C:

```

[edit]
protocols {
bgp {
group int {
type internal;
local-address 10.255.14.176;
neighbor 10.255.14.179;
neighbor 10.255.14.175;
neighbor 10.255.14.177;
neighbor 10.255.14.178;
neighbor 10.255.14.181;
}
}
ospf {
area 0.0.0.0 {
interface lo0.0 {
passive;
}
interface fe-0/3/3.0;
interface fe-0/3/0.0;
}
}
}
}

```

## Configuration for Router D

Router D is the CE router with respect to AS 10023. In a carrier-of-carriers VPN, the CE router must be able to send labels to the carrier provider; this is done with the **labeled-unicast** statement in group **to-isp-red**.

```
[edit]
protocols {
  mpls {
    interface t3-0/0/0.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.175;
      neighbor 10.255.14.179;
      neighbor 10.255.14.176;
      neighbor 10.255.14.177;
      neighbor 10.255.14.178;
      neighbor 10.255.14.181;
    }
    group to-isp-red {
      export internal;
      peer-as 10023;
      neighbor 192.168.197.13 {
        family inet {
          labeled-unicast;
        }
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface fe-0/3/0.0;
    interface t3-0/0/0.0 {
      passive;
    }
  }
}
policy options {
  policy-statement internal {
    term a {
      from protocol [ ospf direct ];
      then accept;
    }
    term b {
      then reject;
    }
  }
}
```

## Configuration for Router E

This configuration sets up the **inet-vpn** IBGP session with Router H and the PE router portion of the VPN with Router D. Because Router D is required to send labels in this example, configure the BGP session with the **labeled-unicast** statement within the virtual routing and forwarding (VRF) table.

```
[edit]
protocols {
  mpls {
    interface t3-0/2/0.0;
    interface at-0/1/0.0;
  }
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.171;
      family inet-vpn {
        any;
      }
      neighbor 10.255.14.173;
    }
  }
  isis {
    interface at-0/1/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface at-0/1/0.0;
  }
}
routing-instances {
  vpn-isp1 {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:21;
    vrf-import vpn-isp1-import;
    vrf-export vpn-isp1-export;
    protocols {
      bgp {
        group to-isp1 {
          peer-as 21;
          neighbor 192.168.197.14 {
            family inet {
              labeled-unicast;
            }
          }
        }
      }
    }
  }
}
```

```
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  community vpn-isp1-comm members target:69:21;
}
```

## Configuration for Router F

Configure Router F to act as a label-swapping router:

```
[edit]
protocols {
  isis {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
  }
}
```

## Configuration for Router G

Configure Router G to act as a label-swapping router:

```
[edit]
protocols {
  isis {
    interface so-0/0/0.0;
    interface so-1/0/0.0;
```



```

        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface so-0/0/0.0;
        interface so-1/0/0.0;
    }
}

```

## Configuration for Router H

Router H acts as the PE router for AS 10023. The configuration that follows is similar to that for Router F:

```

[edit]
protocols {
    mpls {
        interface fe-1/1/0.0;
        interface so-1/0/0.0;
    }
    bgp {
        group pe-pe {
            type internal;
            local-address 10.255.14.173;
            family inet-vpn {
                any;
            }
            neighbor 10.255.14.171;
        }
    }
    isis {
        interface so-1/0/0.0;
        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface so-1/0/0.0;
    }
}
routing-instances {
    vpn-isp1 {
        instance-type vrf;
        interface fe-1/1/0.0;
        route-distinguisher 10.255.14.173:21;
        vrf-import vpn-isp1-import;
        vrf-export vpn-isp1-export;
        protocols {
            bgp {
                group to-isp1 {
                    peer-as 21;
                    neighbor 192.168.197.94 {
                        family inet {
                            labeled-unicast;
                        }
                    }
                }
            }
        }
    }
}

```

```

}
}
}
}
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  community vpn-isp1-comm members target:69:21;
}

```

## Configuration for Router 1

Configure Router I to connect to the basic Internet service customer (Router L):

```
[edit]
protocols {
  mpls {
    interface fe-1/0/1.0;
    interface fe-1/1/3.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.181;
      neighbor 10.255.14.177;
      neighbor 10.255.14.179;
      neighbor 10.255.14.175;
      neighbor 10.255.14.176;
      neighbor 10.255.14.178;
    }
    group to-vpn-green {
      peer-as 1;
    }
  }
}
```

```

        neighbor 192.168.197.198;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-1/0/1.0 {
            passive;
        }
        interface fe-1/1/3.0;
    }
}
}

```

### Configuration for Router J

Configure Router J as a label-swapping router:

```

[edit]
protocols {
    bgp {
        group int {
            type internal;
            local-address 10.255.14.178;
            neighbor 10.255.14.177;
            neighbor 10.255.14.181;
            neighbor 10.255.14.175;
            neighbor 10.255.14.176;
            neighbor 10.255.14.179;
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-1/0/2.0;
        interface fe-1/0/3.0;
    }
}
}

```

### Configuration for Router K

Router K acts as the CE router at the end of the connection to the carrier provider. As in the configuration for Router D, include the **labeled-unicast** statement for the EBGp session:

```

[edit]
protocols {
    mpls {
        interface fe-1/1/2.0;
        interface fe-1/0/2.0;
    }
    bgp {

```

```
group int {
    type internal;
    local-address 10.255.14.177;
    neighbor 10.255.14.181;
    neighbor 10.255.14.178;
    neighbor 10.255.14.175;
    neighbor 10.255.14.176;
    neighbor 10.255.14.179;
}
group to-isp-red {
    export internal;
    peer-as 10023;
    neighbor 192.168.197.93 {
        family inet {
            labeled-unicast;
        }
    }
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-1/0/2.0;
        interface fe-1/1/2.0 {
            passive;
        }
    }
}
}
policy-options {
    policy-statement internal {
        term a {
            from protocol [ ospf direct ];
            then accept;
        }
        term b {
            then reject;
        }
    }
}
}
```

## Configuration for Router L

Configure Router L to act as the end customer for the carrier-of-carriers VPN service:

```
[edit]
protocols {
    bgp {
        group to-routerl {
            export attached;
            peer-as 21;
            neighbor 192.168.197.197;
        }
    }
}
```

```

}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}

```

**Related  
Documentation**

- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)

## Carrier-of-Carriers VPN Example—Customer Provides VPN Service

In this example, the carrier customer *must* run some form of MPLS (Resource Reservation Protocol [RSVP] or LDP) on its network to provide VPN services to the end customer. In the example below, Router B and Router I act as PE routers (or switches), and a functioning MPLS path is required between these routers if they exchange VPN-IPv4 routes.

For configuration information see the following sections:

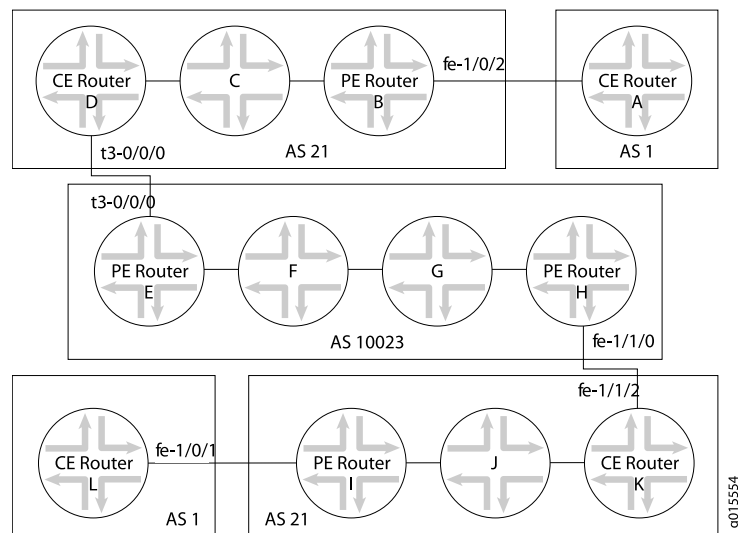
- [Network Topology for Carrier-of-Carriers Service on page 425](#)
- [Configuration for Router A on page 426](#)
- [Configuration for Router B on page 426](#)
- [Configuration for Router C on page 428](#)
- [Configuration for Router D on page 428](#)
- [Configuration for Router E on page 429](#)
- [Configuration for Router F on page 431](#)
- [Configuration for Router G on page 431](#)
- [Configuration for Router H on page 432](#)
- [Configuration for Router I on page 433](#)
- [Configuration for Router J on page 435](#)
- [Configuration for Router K on page 435](#)
- [Configuration for Router L on page 436](#)

### Network Topology for Carrier-of-Carriers Service

A carrier-of-carriers service allows an Internet service provider (ISP) to connect to a transparent outsourced backbone at multiple locations.

[Figure 49 on page 416](#) shows the network topology in this carrier-of-carriers example.

Figure 50: Carrier-of-Carriers VPN Example Network Topology



### Configuration for Router A

In this example, Router A acts as the CE router for the end customer. Configure a default **family inet** BGP session on Router A:

```
[edit]
protocols {
  bgp {
    group to-routerB {
      export attached;
      peer-as 21;
      neighbor 192.168.197.169;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
```

### Configuration for Router B

Because Router B is the PE router for the end customer CE router (Router A), you need to configure a routing instance (**vpna**). Configure the **labeled-unicast** statement on the IBGP session to Router D, and configure **family-inet-vpn** for the IBGP session to the other side of the network with Router I:

```
[edit]
protocols {
  mpls {
    interface fe-1/0/2.0;
    interface fe-1/0/3.0;
  }
}
```

```
bgp {
  group int {
    type internal;
    local-address 10.255.14.179;
    neighbor 10.255.14.175 {
      family inet {
        labeled-unicast {
          resolve-vpn;
        }
      }
    }
  }
  neighbor 10.255.14.181 {
    family inet-vpn {
      any;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface fe-1/0/3.0;
  }
}
ldp {
  interface fe-1/0/3.0;
}
}
routing-instances {
  vpna {
    instance-type vrf;
    interface fe-1/0/2.0;
    route-distinguisher 10.255.14.179:21;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group vpna-06 {
          peer-as 1;
          neighbor 192.168.197.170;
        }
      }
    }
  }
}
policy-options {
  policy-statement vpna-import {
    term a {
      from {
        protocol bgp;
        community vpna-comm;
      }
      then accept;
    }
  }
}
```

```
        term b {
            then reject;
        }
    }
    policy-statement vpna-export {
        term a {
            from protocol bgp;
            then {
                community add vpna-comm;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
    community vpna-comm members target:100:1001;
}
```

## Configuration for Router C

Configure Router C as a label-swapping router within the local AS:

```
[edit]
protocols {
    mpls {
        traffic-engineering bgp-igp;
    }
    ospf {
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fe-0/3/3.0;
            interface fe-0/3/0.0;
        }
    }
    ldp {
        interface fe-0/3/0.0;
        interface fe-0/3/3.0;
    }
}
```

## Configuration for Router D

Router D acts as the CE router for the VPN services provided by the AS 10023 network. In the BGP group configuration for group **int**, which handles traffic to Router B (10.255.14.179), you include the **labeled-unicast** statement. You also need to configure the BGP group **to-isp-red** to send labeled internal routes to the PE router (Router E).

```
[edit]
protocols {
    mpls {
        traffic-engineering bgp-igp;
        interface fe-0/3/0.0;
        interface t3-0/0/0.0;
    }
}
```



```

}
bgp {
  group int {
    type internal;
    local-address 10.255.14.175;
    neighbor 10.255.14.179 {
      family inet {
        labeled-unicast;
      }
    }
  }
  group to-isp-red {
    export internal;
    peer-as 10023;
    neighbor 192.168.197.13 {
      family inet {
        labeled-unicast;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface fe-0/3/0.0;
  }
}
ldp {
  interface fe-0/3/0.0;
}
}
policy-options {
  policy-statement internal {
    term a {
      from protocol [ ospf direct ];
      then accept;
    }
    term b {
      then reject;
    }
  }
}
}

```

## Configuration for Router E

Router E and Router H are PE routers. Configure a PE-router-to-PE-router BGP session to allow VPN-IPv4 routes to pass between these two PE routers. Configure the routing instance on Router E to send labeled routes to the CE router (Router D).

Configure Router E:

```

[edit]
protocols {
  mpls {

```

```
interface t3-0/2/0.0;
interface at-0/1/0.0;
}
bgp {
  group pe-pe {
    type internal;
    local-address 10.255.14.171;
    family inet-vpn {
      any;
    }
    neighbor 10.255.14.173;
  }
}
isis {
  interface at-0/1/0.0;
  interface lo0.0 {
    passive;
  }
}
ldp {
  interface at-0/1/0.0;
}
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  community vpn-isp1-comm members target:69:21;
}
routing-instances {
  vpn-isp1 {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:21;
    vrf-import vpn-isp1-import;
  }
}
```

```

vrf-export vpn-isp1-export;
protocols {
  bgp {
    group to-isp1 {
      peer-as 21;
      neighbor 192.168.197.14 {
        as-override;
        family inet {
          labeled-unicast;
        }
      }
    }
  }
}

```

### Configuration for Router F

Configure Router F to swap labels for routes running through its interfaces:

```

[edit]
protocols {
  isis {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
  }
}

```

### Configuration for Router G

Configure Router G:

```

[edit]
protocols {
  isis {
    interface so-0/0/0.0;
    interface so-1/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-0/0/0.0;
    interface so-1/0/0.0;
  }
}

```

## Configuration for Router H

The configuration for Router H is similar to the configuration for Router E:

```
[edit]
protocols {
  mpls {
    interface fe-1/1/0.0;
    interface so-1/0/0.0;
  }
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.173;
      family inet-vpn {
        any;
      }
      neighbor 10.255.14.171;
    }
  }
  isis {
    interface so-1/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-1/0/0.0;
  }
}
routing-instances {
  vpn-isp1 {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.173:21;
    vrf-import vpn-isp1-import;
    vrf-export vpn-isp1-export;
    protocols {
      bgp {
        group to-isp1 {
          peer-as 21;
          neighbor 192.168.197.94 {
            as-override;
            family inet {
              labeled-unicast;
            }
          }
        }
      }
    }
  }
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
```

```

        from {
            protocol bgp;
            community vpn-isp1-comm;
        }
        then accept;
    }
    term b {
        then reject;
    }
}
policy-statement vpn-isp1-export {
    term a {
        from protocol bgp;
        then {
            community add vpn-isp1-comm;
            accept;
        }
    }
    term b {
        then reject;
    }
}
community vpn-isp1-comm members target:69:21;
}

```

## Configuration for Router I

Router I acts as the PE router for the end customer. The configuration that follows is similar to the configuration for Router B:

```

[edit]
protocols {
    mpls {
        interface fe-1/0/1.0;
        interface fe-1/1/3.0;
    }
    bgp {
        group int {
            type internal;
            local-address 10.255.14.181;
            neighbor 10.255.14.177 {
                family inet {
                    labeled-unicast {
                        resolve-vpn;
                    }
                }
            }
            neighbor 10.255.14.179 {
                family inet-vpn {
                    any;
                }
            }
        }
    }
    ospf {
        area 0.0.0.0 {

```

```
        interface lo0.0 {
            passive;
        }
        interface fe-1/1/3.0;
    }
}
ldp {
    interface fe-1/1/3.0;
}
}
routing-instances {
    vpn {
        instance-type vrf;
        interface fe-1/0/1.0;
        route-distinguisher 10.255.14.181:21;
        vrf-import vpn-import;
        vrf-export vpn-export;
        protocols {
            bgp {
                group vpn-0 {
                    peer-as 1;
                    neighbor 192.168.197.198;
                }
            }
        }
    }
}
}
policy-options {
    policy-statement vpn-import {
        term a {
            from {
                protocol bgp;
                community vpn-comm;
            }
            then accept;
        }
        term b {
            then reject;
        }
    }
    policy-statement vpn-export {
        term a {
            from protocol bgp;
            then {
                community add vpn-comm;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
}
community vpn-comm members target:100:1001;
}
```

## Configuration for Router J

Configure Router J to swap labels for routes running through its interfaces:

```
[edit]
protocols {
  mpls {
    traffic-engineering bgp-igp;
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-1/0/2.0;
      interface fe-1/0/3.0;
    }
  }
  ldp {
    interface fe-1/0/2.0;
    interface fe-1/0/3.0;
  }
}
```

## Configuration for Router K

The configuration for Router K is similar to the configuration for Router D:

```
[edit]
protocols {
  mpls {
    traffic-engineering bgp-igp;
    interface fe-1/1/2.0;
    interface fe-1/0/2.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.177;
      neighbor 10.255.14.181 {
        family inet {
          labeled-unicast;
        }
      }
    }
    group to-isp-red {
      export internal;
      peer-as 10023;
      neighbor 192.168.197.93 {
        family inet {
          labeled-unicast;
        }
      }
    }
  }
}
```

```
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface fe-1/0/2.0;
  }
}
ldp {
  interface fe-1/0/2.0;
}
}
policy-options {
  policy-statement internal {
    term a {
      from protocol [ ospf direct ];
      then accept;
    }
    term b {
      then reject;
    }
  }
}
```

## Configuration for Router L

In this example, Router L is the end customer's CE router. Configure a default family **inet** BGP session on Router L:

```
[edit]
protocols {
  bgp {
    group to-l {
      export attached;
      peer-as 21;
      neighbor 192.168.197.197;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
```

- Related Documentation**
- [MPLS Feature Support on QFX Series and EX4600 Switches](#)
  - [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)

---

## Multiple Instances for LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a core provider PE router to a customer carrier CE router.



Having LDP advertise labels in this manner is especially useful when the carrier customer is a basic ISP and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet at large. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 3 VPN or Layer 2 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Junos Feature Guide* on the product documentation page of the Juniper Networks website, located at <https://www.juniper.net/>.

**Related  
Documentation**

- *MPLS Feature Support on QFX Series and EX4600 Switches*
- [Understanding Interprovider and Carrier-of-Carriers VPNs on page 312](#)



## PART 4

# Configuring Multicast on Layer 3 VPNs

- [Multicast VPN Overview on page 441](#)
- [Configuration Overview on page 475](#)
- [Associating Routing Instances with Sender and Receiver Sites on page 537](#)
- [Providing Resiliency with Redundant Virtual Tunnels on page 567](#)
- [Creating Next Generation MVPN VRF Import and Export Policies on page 587](#)
- [Signaling Provider Tunnels in Next Generation MVPNs on page 593](#)
- [Associating Interfaces with MGBP VPNs on page 599](#)
- [Configuring Draft Rosen VPNs on page 611](#)
- [Configuring Load Balancing on page 623](#)
- [Tracing Traffic and Operations on page 633](#)



## CHAPTER 24

# Multicast VPN Overview

- [Understanding MVPN Concepts and Protocols on page 441](#)
- [MBGP Multicast VPN Sites on page 445](#)
- [Example: Configuring MBGP Multicast VPNs on page 446](#)
- [Configuring Point-to-Multipoint LSPs for an MBGP MVPN on page 464](#)
- [Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471](#)
- [Multicast VPN Terminology on page 473](#)
- [Supported Multicast VPN Standards on page 473](#)

## Understanding MVPN Concepts and Protocols

---

You can configure multicast routing over a network running a Layer 3 VPN that complies with RFC 4364. This section describes this type of network application and includes these topics:

- [Multicast over Layer 3 VPNs Overview on page 441](#)
- [Sending PIM Hello Messages to the PE Routers on page 443](#)
- [Sending PIM Join Messages to the PE Routers on page 444](#)
- [Receiving the Multicast Transmission on page 444](#)

## Multicast over Layer 3 VPNs Overview

In the unicast environment for Layer 3 VPNs, all VPN state information is contained within the PE routers. However, with multicast for Layer 3 VPNs, Protocol Independent Multicast (PIM) adjacencies are established in one of the following ways:

- You can set PIM adjacencies between the CE router and the PE router through a VRF instance at the `[edit routing-instances instance-name protocols pim]` hierarchy level. You must include the `group-address` statement for the provider tunnel, specifying a multicast group. The rendezvous point (RP) listed within the VRF-instance is the VPN customer RP (C-RP).
- You can also set the master PIM instance and the PE's IGP neighbors by configuring statements at the `[edit protocols pim]` hierarchy level. You must add the multicast group specified in the VRF instance to the master PIM instance. The set of master PIM adjacencies throughout the service provider network makes up the forwarding path

that becomes an RP tree rooted at the service provider RP (SP-RP). Therefore, P routers within the provider core must maintain multicast state information for the VPNs.

For this to work properly, you need two types of RP routers for each VPN:

- A C-RP—An RP router located somewhere within the VPN (can be either a service provider router or a customer router).
- An SP-RP—An RP router located within the service provider network.



**NOTE:** A PE router can act as the SP-RP and the C-RP. Moving these multicast configuration tasks to service provider routers helps to simplify the multicast Layer 3 VPN configuration process for customers. However, configuration of both SP-RP and VPN C-RP on the same PE router is not supported.

To configure multicast over a Layer 3 VPN, you must install a Tunnel Services Physical Interface Card (PIC) on the following devices:


- P routers acting as RPs
- PE routers configured to run multicast routing
- CE routers acting as designated routers or as VPN-RPs

For more information about running multicast over Layer 3 VPNs, see the following documents:

- Internet draft draft-rosen-vpn-mcast-02.txt, *Multicast in MPLS/BGP VPNs*
- *Multicast Protocols Feature Guide*

The sections that follow describe the operation of a multicast VPN. [Figure 51 on page 443](#) illustrates the network topology used.

Figure 51: Multicast Topology Overview



## Sending PIM Hello Messages to the PE Routers

The first step in initializing multicast over a Layer 3 VPN is the distribution of a PIM Hello message from a PE router (called PE3 in this section) to all the other PE routers on which PIM is configured.

You configure PIM on the Layer 3 VPN routing instance on the PE3 router. If a Tunnel Services PIC is installed in the routing platform, a multicast interface is created. This interface is used to communicate between the PIM instance within the VRF routing instance and the master PIM instance.

The following occurs when a PIM Hello message is sent to the PE routers:

1. A PIM Hello message is sent from the VRF routing instance over the multicast interface. A generic routing encapsulation (GRE) header is prepended to the PIM Hello message. The header message includes the VPN group address and the loopback address of the PE3 router.
2. A PIM register header is prepended to the Hello message as the packet is looped through the PIM encapsulation interface. This header contains the destination address of the SP-RP and the loopback address of the PE3 router.
3. The packet is sent to the SP-RP.
4. The SP-RP removes the top header from the packet and sends the remaining GRE-encapsulated Hello message to all the PE routers.
5. The master PIM instance on each PE router handles the GRE encapsulated packet. Because the VPN group address is contained in the packet, the master instance removes the GRE header from the packet and sends the Hello message, which contains the proper VPN group address within the VRF routing instance, over the multicast interface.

## Sending PIM Join Messages to the PE Routers

To receive a multicast broadcast from a multicast network, a CE router must send a PIM Join message to the C-RP. The process described in this section refers to [Figure 51 on page 443](#).

The CE5 router needs to receive a multicast broadcast from multicast source 224.1.1.1. To receive the broadcast, it sends a PIM Join message to the C-RP (the PE3 router):

1. The PIM Join message is sent through the multicast interface, and a GRE header is prepended to the message. The GRE header contains the VPN group ID and the loopback address of the PE3 router.
2. The PIM Join message is then sent through the PIM encapsulation interface and a register header is prepended to the packet. The register header contains the IP address of the SP-RP and the loopback address of the PE3 router.
3. The PIM Join message is sent to the SP-RP by means of unicast routing.
4. On the SP-RP, the register header is stripped off (the GRE header remains) and the packet is sent to all the PE routers.
5. The PE2 router receives the packet, and because the link to the C-RP is through the PE2 router, it sends the packet through the multicast interface to remove the GRE header.
6. Finally, the PIM Join message is sent to the C-RP.

## Receiving the Multicast Transmission

The steps that follow outline how a multicast transmission is propagated across the network:

1. The multicast source connected to the CE1 router sends the packet to group 224.1.1.1 (the VPN group address). The packet is encapsulated into a PIM register.
2. Because this packet already includes the PIM header, it is forwarded by means of unicast routing to the C-RP over the Layer 3 VPN.
3. The C-RP removes the packet and sends it out the downstream interfaces (which include the interface back to the CE3 router). The CE3 router also forwards this to the PE3 router.
4. The packet is sent through the multicast interface on the PE2 router; in the process, the GRE header is prepended to the packet.
5. Next, the packet is sent through the PIM encapsulation interface, where the register header is prepended to the data packet.
6. The packet is then forwarded to the SP-RP, which removes the register header, leaves the GRE header intact, and sends the packet to the PE routers.
7. PE routers remove the GRE header and forward the packet to the CE routers that requested the multicast broadcast by sending the PIM Join message.





**NOTE:** PE routers that have not received requests for multicast broadcasts from their connected CE routers still receive packets for the broadcast. These PE routers drop the packets as they are received.

## MBGP Multicast VPN Sites

The main characteristics of MBGP MVPNs are:

- They extend Layer 3 VPN service (RFC 4364) to support IP multicast for Layer 3 VPN service providers.
- They follow the same architecture as specified by RFC 4364 for unicast VPNs. Specifically, BGP is used as the provider edge (PE) router-to-PE router control plane for multicast VPN.
- They eliminate the requirement for the virtual router (VR) model (as specified in Internet draft draft-rosen-vpn-mcast, *Multicast in MPLS/BGP VPNs*) for multicast VPNs and the RFC 4364 model for unicast VPNs.
- They rely on RFC 4364-based unicast with extensions for intra-AS and inter-AS communication.

An MBGP MVPN defines two types of site sets, a sender site set and a receiver site set. These sites have the following properties:

- Hosts within the sender site set can originate multicast traffic for receivers in the receiver site set.
- Receivers outside the receiver site set should not be able to receive this traffic.
- Hosts within the receiver site set can receive multicast traffic originated by any host in the sender site set.
- Hosts within the receiver site set should not be able to receive multicast traffic originated by any host that is not in the sender site set.

A site can be in both the sender site set and the receiver site set, so hosts within such a site can both originate and receive multicast traffic. For example, the sender site set could be the same as the receiver site set, in which case all sites could both originate and receive multicast traffic from one another.

Sites within a given MBGP MVPN might be within the same organization or in different organizations, which means that an MBGP MVPN can be either an intranet or an extranet. A given site can be in more than one MBGP MVPN, so MBGP MVPNs might overlap. Not all sites of a given MBGP MVPN have to be connected to the same service provider, meaning that an MBGP MVPN can span multiple service providers.

Feature parity for the MVPN extranet functionality or overlapping MVPNs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

Another way to look at an MBGP MVPN is to say that an MBGP MVPN is defined by a set of administrative policies. These policies determine both the sender site set and the receiver site set. These policies are established by MBGP MVPN customers, but implemented by service providers using the existing BGP and MPLS VPN infrastructure.

#### Release History Table

Release	Description
11.1R2	Feature parity for the MVPN extranet functionality or overlapping MVPNs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

#### Related Documentation

- [Example: Allowing MBGP MVPN Remote Sources](#)
- [Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN](#)

## Example: Configuring MBGP Multicast VPNs

This example provides a step-by-step procedure to configure multicast services across a multiprotocol BGP (MBGP) Layer 3 virtual private network. (also referred to as next-generation Layer 3 multicast VPNs)

- [Requirements on page 446](#)
- [Overview and Topology on page 447](#)
- [Configuration on page 447](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.2 or later
- Five M Series, T Series, TX Series, or MX Series Juniper routers
- One host system capable of sending multicast traffic and supporting the Internet Group Management Protocol (IGMP)
- One host system capable of receiving multicast traffic and supporting IGMP

Depending on the devices you are using, you might be required to configure static routes to:

- The multicast sender
- The Fast Ethernet interface to which the sender is connected on the multicast receiver
- The multicast receiver
- The Fast Ethernet interface to which the receiver is connected on the multicast sender


## Overview and Topology

This example shows how to configure the following technologies:

- IPv4
- BGP
- OSPF
- RSVP
- MPLS
- PIM sparse mode
- Static RP

The topology of the network is shown in [Figure 52 on page 447](#).

*Figure 52: Multicast Over Layer 3 VPN Example Topology*



## Configuration



**NOTE:** In any configuration session, it is a good practice to periodically verify that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE1** identifies the customer edge 1 (CE1) router
- **PE1** identifies the provider edge 1 (PE1) router
- **P** identifies the provider core (P) router
- **CE2** identifies the customer edge 2 (CE2) router
- **PE2** identifies the provider edge 2 (PE2) router

To configure MBGP multicast VPNs for the network shown in [Figure 52 on page 447](#), perform the following steps:

- [Configuring Interfaces on page 448](#)
- [Configuring OSPF on page 449](#)
- [Configuring BGP on page 450](#)
- [Configuring RSVP on page 451](#)
- [Configuring MPLS on page 452](#)
- [Configuring the VRF Routing Instance on page 453](#)
- [Configuring PIM on page 454](#)
- [Configuring the Provider Tunnel on page 455](#)
- [Configuring the Rendezvous Point on page 455](#)
- [Results on page 456](#)

---

### Configuring Interfaces

---

#### 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*.

1. On each router, configure an IP address on the loopback logical interface 0 (**lo0.0**).

```
[edit interfaces]
```

```
user@CE1# set lo0 unit 0 family inet address 192.168.6.1/32 primary
```

```
user@PE1# set lo0 unit 0 family inet address 192.168.7.1/32 primary
```

```
user@P# set lo0 unit 0 family inet address 192.168.8.1/32 primary
```

```
user@PE2# set lo0 unit 0 family inet address 192.168.9.1/32 primary
```

```
user@CE2# set lo0 unit 0 family inet address 192.168.0.1/32 primary
```

Use the **show interfaces terse** command to verify that the IP address is correct on the loopback logical interface.

2. On the PE and CE routers, configure the IP address and protocol family on the Fast Ethernet interfaces. Specify the **inet** protocol family type.

```
[edit interfaces]
```

```
user@CE1# set fe-1/3/0 unit 0 family inet address 10.10.12.1/24
```

```
user@CE1# set fe-0/1/0 unit 0 family inet address 10.0.67.13/30
```

```
[edit interfaces]
```

```
user@PE1# set fe-0/1/0 unit 0 family inet address 10.0.67.14/30
```

```
[edit interfaces]
```

```
user@PE2# set fe-0/1/0 unit 0 family inet address 10.0.90.13/30
```

```
[edit interfaces]
user@CE2# set fe-0/1/0 unit 0 family inet address 10.0.90.14/30
user@CE2# set fe-1/3/0 unit 0 family inet address 10.10.11.1/24
```

Use the **show interfaces terse** command to verify that the IP address is correct on the Fast Ethernet interfaces.

3. On the PE and P routers, configure the ATM interfaces' VPI and maximum virtual circuits. If the default PIC type is different on directly connected ATM interfaces, configure the PIC type to be the same. Configure the logical interface VCI, protocol family, local IP address, and destination IP address.

```
[edit interfaces]
user@PE1# set at-0/2/0 atm-options pic-type atm1
user@PE1# set at-0/2/0 atm-options vpi 0 maximum-vcs 256
user@PE1# set at-0/2/0 unit 0 vci 0.128
user@PE1# set at-0/2/0 unit 0 family inet address 10.0.78.5/32 destination 10.0.78.6
```

```
[edit interfaces]
user@P# set at-0/2/0 atm-options pic-type atm1
user@P# set at-0/2/0 atm-options vpi 0 maximum-vcs 256
user@P# set at-0/2/0 unit 0 vci 0.128
user@P# set at-0/2/0 unit 0 family inet address 10.0.78.6/32 destination 10.0.78.5
user@P# set at-0/2/1 atm-options pic-type atm1
user@P# set at-0/2/1 atm-options vpi 0 maximum-vcs 256
user@P# set at-0/2/1 unit 0 vci 0.128
user@P# set at-0/2/1 unit 0 family inet address 10.0.89.5/32 destination 10.0.89.6
```

```
[edit interfaces]
user@PE2# set at-0/2/1 atm-options pic-type atm1
user@PE2# set at-0/2/1 atm-options vpi 0 maximum-vcs 256
user@PE2# set at-0/2/1 unit 0 vci 0.128
user@PE2# set at-0/2/1 unit 0 family inet address 10.0.89.6/32 destination 10.0.89.5
```

Use the **show configuration interfaces** command to verify that the ATM interfaces' VPI and maximum VCs are correct and that the logical interface VCI, protocol family, local IP address, and destination IP address are correct.

## Configuring OSPF

### Step-by-Step Procedure

1. On the P and PE routers, configure the provider instance of OSPF. Specify the **lo0.0** and ATM core-facing logical interfaces. The provider instance of OSPF on the PE router forms adjacencies with the OSPF neighbors on the other PE router and Router P.

```
user@PE1# set protocols ospf area 0.0.0.0 interface at-0/2/0.0
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@P# set protocols ospf area 0.0.0.0 interface lo0.0
user@P# set protocols ospf area 0.0.0.0 interface all
user@P# set protocols ospf area 0.0.0.0 interface fxp0 disable
```

```
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.0
user@PE2# set protocols ospf area 0.0.0.0 interface at-0/2/1.0
```

Use the **show ospf interfaces** command to verify that the **lo0.0** and ATM core-facing logical interfaces are configured for OSPF.

2. On the CE routers, configure the customer instance of OSPF. Specify the loopback and Fast Ethernet logical interfaces. The customer instance of OSPF on the CE routers form adjacencies with the neighbors within the VPN routing instance of OSPF on the PE routers.

```
user@CE1# set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
user@CE1# set protocols ospf area 0.0.0.0 interface fe-1/3/0.0
user@CE1# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@CE2# set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
user@CE2# set protocols ospf area 0.0.0.0 interface fe-1/3/0.0
user@CE2# set protocols ospf area 0.0.0.0 interface lo0.0
```

Use the **show ospf interfaces** command to verify that the correct loopback and Fast Ethernet logical interfaces have been added to the OSPF protocol.

3. On the P and PE routers, configure OSPF traffic engineering support for the provider instance of OSPF.

The **shortcuts** statement enables the master instance of OSPF to use a label-switched path as the next hop.

```
user@PE1# set protocols ospf traffic-engineering shortcuts
```

```
user@P# set protocols ospf traffic-engineering shortcuts
```

```
user@PE2# set protocols ospf traffic-engineering shortcuts
```

Use the **show ospf overview** or **show configuration protocols ospf** command to verify that traffic engineering support is enabled.

## Configuring BGP

### Step-by-Step Procedure

1. On Router P, configure BGP for the VPN. The local address is the local **lo0.0** address. The neighbor addresses are the PE routers' **lo0.0** addresses.

The **unicast** statement enables the router to use BGP to advertise network layer reachability information (NLRI). The **signaling** statement enables the router to use BGP as the signaling protocol for the VPN.

```
user@P# set protocols bgp group group-mvpn type internal
user@P# set protocols bgp group group-mvpn local-address 192.168.8.1
user@P# set protocols bgp group group-mvpn family inet unicast
user@P# set protocols bgp group group-mvpn family inet-mvpn signaling
user@P# set protocols bgp group group-mvpn neighbor 192.168.9.1
user@P# set protocols bgp group group-mvpn neighbor 192.168.7.1
```

Use the **show configuration protocols bgp** command to verify that the router has been configured to use BGP to advertise NLRI.

2. On the PE and P routers, configure the BGP local autonomous system number.

```
user@PE1# set routing-options autonomous-system 0.65010
```

```
user@P# set routing-options autonomous-system 0.65010
```

```
user@PE2# set routing-options autonomous-system 0.65010
```

Use the **show configuration routing-options** command to verify that the BGP local autonomous system number is correct.

3. On the PE routers, configure BGP for the VPN. Configure the local address as the local **lo0.0** address. The neighbor addresses are the **lo0.0** addresses of Router P and the other PE router, PE2.

```
user@PE1# set protocols bgp group group-mvpn type internal
user@PE1# set protocols bgp group group-mvpn local-address 192.168.7.1
user@PE1# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE1# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.9.1
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.8.1
```

```
user@PE2# set protocols bgp group group-mvpn type internal
user@PE2# set protocols bgp group group-mvpn local-address 192.168.9.1
user@PE2# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE2# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.7.1
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.8.1
```

Use the **show bgp group** command to verify that the BGP configuration is correct.

4. On the PE routers, configure a policy to export the BGP routes into OSPF.

```
user@PE1# set policy-options policy-statement bgp-to-ospf from protocol bgp
user@PE1# set policy-options policy-statement bgp-to-ospf then accept
```

```
user@PE2# set policy-options policy-statement bgp-to-ospf from protocol bgp
user@PE2# set policy-options policy-statement bgp-to-ospf then accept
```

Use the **show policy bgp-to-ospf** command to verify that the policy is correct.

## Configuring RSVP

### Step-by-Step Procedure

1. On the PE routers, enable RSVP on the interfaces that participate in the LSP. Configure the Fast Ethernet and ATM logical interfaces.

```
user@PE1# set protocols rsvp interface fe-0/1/0.0
user@PE1# set protocols rsvp interface at-0/2/0.0
```

```
user@PE2# set protocols rsvp interface fe-0/1/0.0
user@PE2# set protocols rsvp interface at-0/2/1.0
```

2. On Router P, enable RSVP on the interfaces that participate in the LSP. Configure the ATM logical interfaces.

```
user@P# set protocols rsvp interface at-0/2/0.0
user@P# set protocols rsvp interface at-0/2/1.0
```

Use the **show configuration protocols rsvp** command to verify that the RSVP configuration is correct.

### Configuring MPLS

#### Step-by-Step Procedure

1. On the PE routers, configure an MPLS LSP to the PE router that is the LSP egress point. Specify the IP address of the **lo0.0** interface on the router at the other end of the LSP. Configure MPLS on the ATM, Fast Ethernet, and **lo0.0** interfaces.

To help identify each LSP when troubleshooting, configure a different LSP name on each PE router. In this example, we use the name **to-pe2** as the name for the LSP configured on PE1 and **to-pe1** as the name for the LSP configured on PE2.

```
user@PE1# set protocols mpls label-switched-path to-pe2 to 192.168.9.1
user@PE1# set protocols mpls interface fe-0/1/0.0
user@PE1# set protocols mpls interface at-0/2/0.0
user@PE1# set protocols mpls interface lo0.0
```

```
user@PE2# set protocols mpls label-switched-path to-pe1 to 192.168.7.1
user@PE2# set protocols mpls interface fe-0/1/0.0
user@PE2# set protocols mpls interface at-0/2/1.0
user@PE2# set protocols mpls interface lo0.0
```

Use the **show configuration protocols mpls** and **show route label-switched-path to-pe1** commands to verify that the MPLS and LSP configuration is correct.

After the configuration is committed, use the **show mpls lsp name to-pe1** and **show mpls lsp name to-pe2** commands to verify that the LSP is operational.

2. On Router P, enable MPLS. Specify the ATM interfaces connected to the PE routers.

```
user@P# set protocols mpls interface at-0/2/0.0
user@P# set protocols mpls interface at-0/2/1.0
```

Use the **show mpls interface** command to verify that MPLS is enabled on the ATM interfaces.

3. On the PE and P routers, configure the protocol family on the ATM interfaces associated with the LSP. Specify the **mpls** protocol family type.

```
user@PE1# set interfaces at-0/2/0 unit 0 family mpls
```

```
user@P# set interfaces at-0/2/0 unit 0 family mpls
user@P# set interfaces at-0/2/1 unit 0 family mpls
```



```
user@PE2# set interfaces at-0/2/1 unit 0 family mpls
```

Use the **show mpls interface** command to verify that the MPLS protocol family is enabled on the ATM interfaces associated with the LSP.

### Configuring the VRF Routing Instance

#### Step-by-Step Procedure

1. On the PE routers, configure a routing instance for the VPN and specify the **vrf** instance type. Add the Fast Ethernet and **lo0.1** customer-facing interfaces. Configure the VPN instance of OSPF and include the BGP-to-OSPF export policy.

```
user@PE1# set routing-instances vpn-a instance-type vrf
user@PE1# set routing-instances vpn-a interface lo0.1
user@PE1# set routing-instances vpn-a interface fe-0/1/0.0
user@PE1# set routing-instances vpn-a protocols ospf export bgp-to-ospf
user@PE1# set routing-instances vpn-a protocols ospf area 0.0.0.0 interface all
```

```
user@PE2# set routing-instances vpn-a instance-type vrf
user@PE2# set routing-instances vpn-a interface lo0.1
user@PE2# set routing-instances vpn-a interface fe-0/1/0.0
user@PE2# set routing-instances vpn-a protocols ospf export bgp-to-ospf
user@PE2# set routing-instances vpn-a protocols ospf area 0.0.0.0 interface all
```

Use the **show configuration routing-instances vpn-a** command to verify that the routing instance configuration is correct.

2. On the PE routers, configure a route distinguisher for the routing instance. A route distinguisher allows the router to distinguish between two identical IP prefixes used as VPN routes. Configure a different route distinguisher on each PE router. This example uses 65010:1 on PE1 and 65010:2 on PE2.

```
user@PE1# set routing-instances vpn-a route-distinguisher 65010:1
```

```
user@PE2# set routing-instances vpn-a route-distinguisher 65010:2
```

Use the **show configuration routing-instances vpn-a** command to verify that the route distinguisher is correct.

3. On the PE routers, configure default VRF import and export policies. Based on this configuration, BGP automatically generates local routes corresponding to the route target referenced in the VRF import policies. This example uses 2:1 as the route target.



**NOTE:** You must configure the same route target on each PE router for a given VPN routing instance.

```
user@PE1# set routing-instances vpn-a vrf-target target:2:1
```

```
user@PE2# set routing-instances vpn-a vrf-target target:2:1
```

Use the **show configuration routing-instances vpn-a** command to verify that the route target is correct.

4. On the PE routers, configure the VPN routing instance for multicast support.

```
user@PE1# set routing-instances vpn-a protocols mvpn
```

```
user@PE2# set routing-instances vpn-a protocols mvpn
```

Use the **show configuration routing-instance vpn-a** command to verify that the VPN routing instance has been configured for multicast support.

5. On the PE routers, configure an IP address on loopback logical interface 1 (**lo0.1**) used in the customer routing instance VPN.

```
user@PE1# set interfaces lo0 unit 1 family inet address 10.10.47.101/32
```

```
user@PE2# set interfaces lo0 unit 1 family inet address 10.10.47.100/32
```

Use the **show interfaces terse** command to verify that the IP address on the loopback interface is correct.

---

### Configuring PIM

---

#### Step-by-Step Procedure

1. On the PE routers, enable PIM. Configure the **lo0.1** and the customer-facing Fast Ethernet interface. Specify the mode as **sparse** and the version as **2**.

```
user@PE1# set routing-instances vpn-a protocols pim interface lo0.1 mode sparse
user@PE1# set routing-instances vpn-a protocols pim interface lo0.1 version 2
user@PE1# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 mode sparse
user@PE1# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 version 2

user@PE2# set routing-instances vpn-a protocols pim interface lo0.1 mode sparse
user@PE2# set routing-instances vpn-a protocols pim interface lo0.1 version 2
user@PE2# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 mode sparse
user@PE2# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 version 2
```

Use the **show pim interfaces instance vpn-a** command to verify that PIM sparse-mode is enabled on the **lo0.1** interface and the customer-facing Fast Ethernet interface.
2. On the CE routers, enable PIM. In this example, we configure all interfaces. Specify the mode as **sparse** and the version as **2**.

```
user@CE1# set protocols pim interface all

user@CE2# set protocols pim interface all mode sparse
user@CE2# set protocols pim interface all version 2
```

Use the **show pim interfaces** command to verify that PIM sparse mode is enabled on all interfaces.

### Configuring the Provider Tunnel

#### Step-by-Step Procedure

1. On Router PE1, configure the provider tunnel. Specify the multicast address to be used.  
  
The **provider-tunnel** statement instructs the router to send multicast traffic across a tunnel.  
  

```
user@PE1# set routing-instances vpn-a provider-tunnel rsvp-te label-switched-path-template default-template
```

  
Use the **show configuration routing-instance vpn-a** command to verify that the provider tunnel is configured to use the default LSP template.
2. On Router PE2, configure the provider tunnel. Specify the multicast address to be used.  
  

```
user@PE2# set routing-instances vpn-a provider-tunnel rsvp-te label-switched-path-template default-template
```

  
Use the **show configuration routing-instance vpn-a** command to verify that the provider tunnel is configured to use the default LSP template.

### Configuring the Rendezvous Point

#### Step-by-Step Procedure

1. Configure Router PE1 to be the rendezvous point. Specify the **lo0.1** address of Router PE1. Specify the multicast address to be used.  
  

```
user@PE1# set routing-instances vpn-a protocols pim rp local address 10.10.47.101
user@PE1# set routing-instances vpn-a protocols pim rp local group-ranges 224.1.1.1/32
```

  
Use the **show pim rps instance vpn-a** command to verify that the correct local IP address is configured for the RP.
2. On Router PE2, configure the static rendezvous point. Specify the **lo0.1** address of Router PE1.  
  

```
user@PE2# set routing-instances vpn-a protocols pim rp static address 10.10.47.101
```

  
Use the **show pim rps instance vpn-a** command to verify that the correct static IP address is configured for the RP.
3. On the CE routers, configure the static rendezvous point. Specify the **lo0.1** address of Router PE1.  
  

```
user@CE1# set protocols pim rp static address 10.10.47.101 version 2
```

```
user@CE2# set protocols pim rp static address 10.10.47.101 version 2
```

Use the **show pim rps** command to verify that the correct static IP address is configured for the RP.

4. Use the **commit check** command to verify that the configuration can be successfully committed. If the configuration passes the check, commit the configuration.
5. Start the multicast sender device connected to CE1.
6. Start the multicast receiver device connected to CE2.
7. Verify that the receiver is receiving the multicast stream.
8. Use **show** commands to verify the routing, VPN, and multicast operation.

---

## Results

The configuration and verification parts of this example have been completed. The following section is for your reference.

The relevant sample configuration for Router CE1 follows.

```
Router CE1 interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.168.6.1/32 {
          primary;
        }
      }
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.67.13/30;
      }
    }
  }
  fe-1/3/0 {
    unit 0 {
      family inet {
        address 10.10.12.1/24;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface fe-0/1/0.0;
    }
  }
}
```

```

        interface lo0.0;
        interface fe-1/3/0.0;
    }
}
pim {
    rp {
        static {
            address 10.10.47.101 {
                version 2;
            }
        }
    }
    interface all;
}
}

```

The relevant sample configuration for Router PE1 follows.

```

Router PE1  interfaces {
              lo0 {
                unit 0 {
                  family inet {
                    address 192.168.7.1/32 {
                      primary;
                    }
                  }
                }
              }
              fe-0/1/0 {
                unit 0 {
                  family inet {
                    address 10.0.67.14/30;
                  }
                }
              }
              at-0/2/0 {
                atm-options {
                  pic-type atm1;
                  vpi 0 {
                    maximum-vcs 256;
                  }
                }
                unit 0 {
                  vci 0.128;
                  family inet {
                    address 10.0.78.5/32 {
                      destination 10.0.78.6;
                    }
                  }
                }
                family mpls;
              }
            }
            lo0 {
              unit 1 {
                family inet {
                  address 10.10.47.101/32;
                }
              }
            }
          }

```

```
    }
  }
}
routing-options {
  autonomous-system 0.65010;
}
protocols {
  rsvp {
    interface fe-0/1/0.0;
    interface at-0/2/0.0;
  }
  mpls {
    label-switched-path to-pe2 {
      to 192.168.9.1;
    }
    interface fe-0/1/0.0;
    interface at-0/2/0.0;
    interface lo0.0;
  }
  bgp {
    group group-mvpn {
      type internal;
      local-address 192.168.7.1;
      family inet-vpn {
        unicast;
      }
      family inet-mvpn {
        signaling;
      }
      neighbor 192.168.9.1;
      neighbor 192.168.8.1;
    }
  }
  ospf {
    traffic-engineering {
      shortcuts;
    }
    area 0.0.0.0 {
      interface at-0/2/0.0;
      interface lo0.0;
    }
  }
}
policy-options {
  policy-statement bgp-to-ospf {
    from protocol bgp;
    then accept;
  }
}
routing-instances {
  vpn-a {
    instance-type vrf;
    interface lo0.1;
    interface fe-0/1/0.0;
    route-distinguisher 65010:1;
  }
}
```

```

provider-tunnel {
  rsvp-te {
    label-switched-path-template {
      default-template;
    }
  }
}
vrf-target target:2:1;
protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.0 {
      interface all;
    }
  }
  pim {
    rp {
      local {
        address 10.10.47.101;
        group-ranges {
          224.1.1.1/32;
        }
      }
    }
    interface lo0.1 {
      mode sparse;
      version 2;
    }
    interface fe-0/1/0.0 {
      mode sparse;
      version 2;
    }
  }
  mvpn;
}
}
}

```

The relevant sample configuration for Router P follows.

```

Router P  interfaces {
          lo0 {
            unit 0 {
              family inet {
                address 192.168.8.1/32 {
                  primary;
                }
              }
            }
          }
          at-0/2/0 {
            atm-options {
              pic-type atm1;
              vpi 0 {
                maximum-vcs 256;
              }
            }
          }
        }

```

```
}
unit 0 {
  vci 0.128;
  family inet {
    address 10.0.78.6/32 {
      destination 10.0.78.5;
    }
  }
  family mpls;
}
}
at-0/2/1 {
  atm-options {
    pic-type atm1;
    vpi 0 {
      maximum-vcs 256;
    }
  }
  unit 0 {
    vci 0.128;
    family inet {
      address 10.0.89.5/32 {
        destination 10.0.89.6;
      }
    }
    family mpls;
  }
}
}
routing-options {
  autonomous-system 0.65010;
}
protocols {
  rsvp {
    interface at-0/2/0.0;
    interface at-0/2/1.0;
  }
  mpls {
    interface at-0/2/0.0;
    interface at-0/2/1.0;
  }
  bgp {
    group group-mvpn {
      type internal;
      local-address 192.168.8.1;
      family inet {
        unicast;
      }
      family inet-mvpn {
        signaling;
      }
      neighbor 192.168.9.1;
      neighbor 192.168.7.1;
    }
  }
  ospf {
```



```

        traffic-engineering {
            shortcuts;
        }
        area 0.0.0.0 {
            interface lo0.0;
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}

```

The relevant sample configuration for Router PE2 follows.

```

Router PE2 interfaces {
    lo0 {
        unit 0 {
            family inet {
                address 192.168.9.1/32 {
                    primary;
                }
            }
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.13/30;
            }
        }
    }
    at-0/2/1 {
        atm-options {
            pic-type atm1;
            vpi 0 {
                maximum-vcs 256;
            }
        }
        unit 0 {
            vci 0.128;
            family inet {
                address 10.0.89.6/32 {
                    destination 10.0.89.5;
                }
            }
            family mpls;
        }
    }
    lo0 {
        unit 1 {
            family inet {
                address 10.10.47.100/32;
            }
        }
    }
}

```

```
}
routing-options {
    autonomous-system 0.65010;
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface at-0/2/1.0;
    }
    mpls {
        label-switched-path to-pe1 {
            to 192.168.7.1;
        }
        interface lo0.0;
        interface fe-0/1/0.0;
        interface at-0/2/1.0;
    }
    bgp {
        group group-mvpn {
            type internal;
            local-address 192.168.9.1;
            family inet-vpn {
                unicast;
            }
            family inet-mvpn {
                signaling;
            }
            neighbor 192.168.7.1;
            neighbor 192.168.8.1;
        }
    }
    ospf {
        traffic-engineering {
            shortcuts;
        }
        area 0.0.0.0 {
            interface lo0.0;
            interface at-0/2/1.0;
        }
    }
}
policy-options {
    policy-statement bgp-to-ospf {
        from protocol bgp;
        then accept;
    }
}
routing-instances {
    vpn-a {
        instance-type vrf;
        interface fe-0/1/0.0;
        interface lo0.1;
        route-distinguisher 65010:2;
        provider-tunnel {
            rsvp-te {
                label-switched-path-template {
```

```

        default-template;
    }
}
vrf-target target:2:1;
protocols {
    ospf {
        export bgp-to-ospf;
        area 0.0.0.0 {
            interface all;
        }
    }
    pim {
        rp {
            static {
                address 10.10.47.101;
            }
        }
        interface fe-0/1/0.0 {
            mode sparse;
            version 2;
        }
        interface lo0.1 {
            mode sparse;
            version 2;
        }
    }
}
mvpn;
}
}

```

The relevant sample configuration for Router CE2 follows.

```

Router CE2 interfaces {
    lo0 {
        unit 0 {
            family inet {
                address 192.168.0.1/32 {
                    primary;
                }
            }
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.14/30;
            }
        }
    }
    fe-1/3/0 {
        unit 0 {
            family inet {
                address 10.10.11.1/24;
            }
        }
    }
}

```

```

        family inet6 {
            address fe80::205:85ff:fe88:ccdb/64;
        }
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface lo0.0;
            interface fe-1/3/0.0;
        }
    }
    pim {
        rp {
            static {
                address 10.10.47.101 {
                    version 2;
                }
            }
        }
        interface all {
            mode sparse;
            version 2;
        }
    }
}
}

```

[Related  
Documentation](#)

## Configuring Point-to-Multipoint LSPs for an MBGP MVPN

The Junos OS supports point-to-multipoint label-switched paths (LSPs) for MBGP MVPNs. Point-to-multipoint LSPs for multicast VPNs are supported for intra-autonomous system (AS) environments (within an AS), but are not supported for inter-AS environments (between autonomous systems). A point-to-multipoint LSP is an RSVP-signaled LSP with a single source and multiple destinations.

You can configure point-to-multipoint LSPs for MBGP MVPNs as follows:

- Static point-to-multipoint LSPs—Configure static point-to-multipoint LSPs using the standard MPLS LSP statements specified at the **[edit protocols mpls]** hierarchy level. You manually configure each of the leaf nodes for the point-to-multipoint LSP.
- Dynamic point-to-multipoint LSPs using the default template—Configuring dynamic point-to-multipoint LSPs using the **default-template** option causes the leaf nodes to be discovered automatically. The leaf nodes are discovered through BGP intra-AS automatic discovery. The **default-template** option allows you to minimize the amount of configuration needed. However, it does not allow you to configure any of the standard MPLS options.

- Dynamic point-to-multipoint LSPs using a user-configured template—Configuring dynamic point-to-multipoint LSPs using a user-configured template also causes the leaf nodes to be discovered automatically. By creating your own template for the point-to-multipoint LSPs, all of the standard MPLS features (such as bandwidth allocation and traffic engineering) can be configured.

Be aware of the following properties for the egress PE router in a point-to-multipoint LSP configured for a multicast VPN:

- Penultimate hop-popping is not used by point-to-multipoint LSPs for multicast VPNs. Only ultimate hop-popping is used.
- You must configure either the **vrf-table-label** statement or a virtual loopback tunnel interface on the egress PE router.
- If you configure the **vrf-table-label** statement on the egress PE router, and the egress PE router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends two copies of each packet over the link to the egress PE router.
- If you configure the **vrf-table-label** statement on the egress PE router, and the egress PE router is not a transit router for the point-to-multipoint LSP, the penultimate hop router can send just one copy of each packet over the link to the egress PE router.
- If you configure a virtual loopback tunnel interface on the egress PE router, and the egress PE router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends just one copy of each packet over the link to the egress PE router. A virtual loopback tunnel interface can perform two lookups on an incoming packet, one for the multicast MPLS lookup and one for the IP lookup.



**NOTE:** Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.

The following sections describe how to configure point-to-multipoint LSPs for MBGP MVPNs:

- [Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 465](#)
- [Configuring Selective Provider Tunnels for an MBGP MVPN on page 466](#)

## Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN

You can configure LDP-signaled or RSVP-signaled inclusive point-to-multipoint LSPs for MBGP MVPNs. Aggregation is not supported, so you need to configure an inclusive point-to-multipoint LSP for each sender PE router in each multicast VPN routing instance. The sender PE router is in the sender site set of the MBGP MVPN.

To configure a static RSVP-signaled inclusive point-to-multipoint LSP, include the **static-lsp** statement:

```
static-lsp lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]

To configure dynamic inclusive point-to-multipoint LSPs, include the **label-switched-path-template** statement:

```
label-switched-path-template (Multicast) {
  (default-template | lsp-template-name);
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]

You can configure either the **default-template** option or manually configure a point-to-multipoint LSP template and specify the template name.

## Configuring Selective Provider Tunnels for an MBGP MVPN

You can configure LDP-signaled or RSVP-signaled selective point-to-multipoint LSPs (also referred to as selective provider tunnels) for MBGP MVPNs. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the multicast VPNs, helping to minimize flooding in the service provider's network.

As with inclusive point-to-multipoint LSPs, you can configure both dynamic and static selective tunnels for the multicast VPN.

To configure selective point-to-multipoint provider tunnels, include the **selective** statement:

```
selective {
  group multicast--prefix/prefix-length {
    source ip--prefix/prefix-length {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbps;
  }
  wildcard-source {
    ldp-p2mp;
    pim-ssm {
      group-range multicast-prefix;
    }
  }
}
```

```

    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kpbs;
}
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
    }
    threshold-rate number;
}
}
wildcard-group-inet6 {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
    }
    threshold-rate number;
}
}
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]

The following sections describe how to configure selective point-to-multipoint LSPs for MBGP MVPNs:

- [Configuring the Multicast Group Address for an MBGP MVPN on page 468](#)
- [Configuring the Multicast Source Address for an MBGP MVPN on page 468](#)

- [Configuring Static Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 468](#)
- [Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 469](#)
- [Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 469](#)
- [Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 470](#)

---

### Configuring the Multicast Group Address for an MBGP MVPN

To configure a point-to-multipoint LSP for an MBGP MVPN, you need to specify a multicast group address by including the **group** statement:

```
group address { ... }
```

You can include this statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective]

The address must be a valid multicast group address. Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255).

---

### Configuring the Multicast Source Address for an MBGP MVPN

To configure a point-to-multipoint LSP for an MBGP MVPN, specify a multicast source address by including the **source** statement:

```
source address { ... }
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective group address]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address]

---

### Configuring Static Selective Point-to-Multipoint LSPs for an MBGP MVPN

You can configure a static selective point-to-multipoint LSP for an MBGP MVPN. You need to configure a static LSP using the standard MPLS LSP statements at the [edit protocols mpls] hierarchy level. You then include the static LSP in your selective point-to-multipoint LSP configuration by using the **static-lsp** statement. Once this functionality is enabled on the source PE router, the static point-to-multipoint LSP is created based on your configuration.

To configure a static selective point-to-multipoint LSP, include the **rsvp-te** and the **static-lsp** statements:

```
rsvp-te static-lsp lsp-name;
```



You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]

### Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

You can configure a dynamic selective point-to-multipoint LSP for an MBGP MVPN. The leaf nodes for a dynamic point-to-multipoint LSP can be automatically discovered using leaf automatic discovery routes. Selective provider multicast service interface (S-PMSI) automatic discovery routes are also supported.

To configure a dynamic selective point-to-multipoint provider tunnel, include the **rsvp-te** and **label-switched-path-template** statements:

```
rsvp-te label-switched-path-template {
  (default-template | lsp-template-name);
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]

The **label-switched-path-template** statement includes the following options:

- **default-template**—Specify that point-to-multipoint LSPs are generated dynamically based on the default template. No user configuration is required for the LSPs. However, the automatically generated LSPs include none of the common LSP features, such as bandwidth allocation and traffic engineering.
- ***lsp-template-name***—Specify the name of an LSP template to be used for the point-to-multipoint LSP. You need to configure the LSP template to be used as a basis for the point-to-multipoint LSPs. You can configure any of the common LSP features for this template.

### Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

To configure a selective point-to-multipoint LSP dynamically, you need to specify the data threshold (in kilobits per second) required before a new tunnel is created using the **threshold-rate** statement:

```
threshold-rate number;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]

### Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

---

To configure a limit on the number of tunnels that can be generated for a dynamic point-to-multipoint LSP, include the **tunnel-limit** statement:

**tunnel-limit** *number*;

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective]

#### Related Documentation

- *Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs*

## Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview

Junos OS supports point-to-multipoint (P2MP) label-switched paths (LSPs) for BGP MVPNs. BGP MVPN supports non-segmented intra-autonomous systems (ASs) and segmented inter-autonomous systems (ASs).

In order to connect PE routers that are in different areas but in the same AS and require P2MP connectivity, Junos OS allows you to segment the P2MP LSPs at the area boundary as described in Internet draft *draft-ietf-mpls-seamless-mcast-14.txt*. You can use non-segmented LSPs for low-rate multicast flows, and segmented LSPs for high-rate flows. A segmented P2MP LSP within an AS consists of the following segments:

- Ingress area segment — The ingress area segment is rooted at a PE router or autonomous system boundary router (ASBR). The leaves of this segment are PEs, ASBRs, or area border routers (ABRs).
- Backbone area segment — The backbone area segment is rooted at an ABR that is connected to the ingress area/ingress ABR.
- Egress area segment — The egress area segment is rooted at an ABR in the egress area or egress ABR.



**NOTE:** These areas can be IGP areas or areas based on BGP peer groups, where ABR can be a region border router (RBR). In either case, the transit ABRs/RBRs should be configured on the BGP route reflector (RR).

Each of the intra-area segments can be carried over provider tunnels such as P2MP RSVP-TE LSP, P2MP mLDP LSP, or ingress replication.

Segmentation of inter-area P2MP LSP occurs when the S-PMSI autodiscovery (AD) routes are advertised. This triggers the inclusion of a new BGP extended community or inter-area P2MP segmented next-hop extended community. The segmented inter-area P2MP LSP can be separated into the following three different roles:

- Ingress PE or ASBR — Ingress PE router originates S-PMSI A-D routes. If inter-region segmentation is required, then the PE router generates the S-PMSI A-D routes carrying the inter-area P2MP segmented next-hop router (S-NH) community. The inter-region segmentation can be added for any selective tunnel. The segmentation can happen based on the threshold or fan-out attributes. If the threshold is configured for a selective tunnel, then MVPN starts migrating the flow to a segmented S-PMSI on reaching the threshold rate value. The threshold attribute applies to RSVP, LDP, and IR tunnels. You can trigger the segmentation based on the fan-out attribute, which is the number of leaves. Once the number of leaf A-D routes exceeds the fan-out value, the traffic flow is moved to segmented S-PMSI. The fan-out attribute for LDP tunnels is not applicable at the ingress PE router.  
If the S-PMSI with ingress replication has configured only the threshold, then the threshold is used to trigger the migration to segmented LSP. If fan-out is also set, then the migration is triggered when the traffic rate multiplied by the number of leaf A-D routes exceeds the threshold value. The segmented threshold and fan-out values are

checked based on the existing data threshold checking interval, which by default is every 60 seconds. This prevents the flow from getting migrated too frequently.

- **Transit ABRs** — When the transit ABR (either ingress ABR or egress ABR) receives an S-PMSI A-D route with the segmentation of inter-region configured, the ABR checks if the S-PMSI is carrying a S-NH extended community attribute. If the S-NH attribute is present in the incoming S-PMSI, then the ABR checks for the tunnel-type to be carried by the S-PMSI. The ABR then generates the tunnel-type across the backbone area or the egress area .



**NOTE:** An ABR can set a template to define the provider tunnel type in each region or BGP group. The tunnel type in each region can be incoming, ingress-replication, LDP-P2MP, or RSVP-TE.

If the tunnel type is incoming, then it indicates that the tunnel type across the ABR remains the same. If the tunnel type is different across the ABR, then the transit ABR modifies the S-PMSI tunnel attribute and the S-NH attribute to its router-id and re-advertises the route to its BGP peers. If no template is configured on the ABR then the ABR simply reflects the incoming S-PMSI routes without changing any of the attributes to its BGP peers.

- **Egress PE or ASBR** — Egress PE routers or ASBRs learn the upstream node from the segmented next-hop extended community carried in the received S-PMSI A-D routes and responds with the leaf A-D routes carrying the upstream node IP address in the route target extended community (EC).

You can configure the BGP policy to accept or reject the S-PMSI A-D routes carrying the inter-area P2MP segmented next-hop community.

#### Related Documentation

- [Example: Configuring Segmented Inter-Area P2MP LSP on page 478](#)
- [Configuring Segmented Inter-Area P2MP LSP on page 476](#)
- [all-regions on page 1042](#)
- [inter-region on page 1082](#)
- [inter-region-segmented on page 1083](#)
- [inter-region-template on page 1085](#)
- [region on page 1122](#)
- [template on page 1141](#)

## Multicast VPN Terminology

### I

**Inclusive tree** A single multicast distribution tree in the backbone that carries all the multicast traffic from a specified set of one or more multicast VPNs. An inclusive tree that carries the traffic of more than one multicast VPN is an aggregate inclusive tree. An inclusive tree contains as its members all the PE routers that attach to the receiver sites of any of the multicast VPNs using the tree.

### S

**Selective tree** A single multicast distribution tree in the backbone that carries traffic belonging only to a specified set of one or more multicast groups, from one or more multicast VPNs. An aggregate selective tree carries traffic for multicast groups that belong to different multicast VPNs. By default, traffic from most multicast groups could be carried by an inclusive tree, whereas traffic from high-bandwidth groups should be carried by a selective tree.

## Supported Multicast VPN Standards

Junos OS substantially supports the following RFCs and Internet draft, which define standards for multicast virtual private networks (VPNs).

- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*
- RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPN*
- RFC 6625, *Wildcards in Multicast VPN Auto-Discovery Routes*
- Internet draft draft-morin-l3vpn-mvpn-fast-failover-06.txt, *Multicast VPN Fast Upstream Failover*
- Internet draft draft-raggarwa-l3vpn-bgp-mvpn-extranet-08.txt, *Extranet in BGP Multicast VPN (MVPN)*

### Related Documentation

- [Supported Carrier-of-Carriers and Interprovider VPN Standards on page 313](#)
- [Supported VPWS Standards](#)
- [Supported Layer 2 VPN Standards](#)
- [Supported Layer 3 VPN Standards on page 130](#)
- [Supported VPLS Standards](#)
- [Supported MPLS Standards](#)
- [Supported Standards for BGP](#)
- [Accessing Standards Documents on the Internet](#)



## CHAPTER 25

# Configuration Overview

- [Configuring Segmented Inter-Area P2MP LSP on page 476](#)
- [Example: Configuring Segmented Inter-Area P2MP LSP on page 478](#)

## Configuring Segmented Inter-Area P2MP LSP

In order to connect PE routers that are in different areas but in the same AS and that require P2MP connectivity, Junos OS allows you to segment the P2MP LSPs at the area boundary as described in Internet draft *draft-ietf-mpls-seamless-mcast-14.txt*.

To configure segmented inter-area P2MP LSPs at the ingress area segment, the backbone area segment, and the egress area segment, you must do the following:

1. Configure inter-region-segmented for group, wildcard-group-inet, or wildcard-group-inet6 of selective tunnel.
  - Configure inter-region-segmented fan-out and threshold values for a multicast source or wildcard source belonging to group.
    - Specify fan-out and threshold values for a multicast source.
 

```
[edit routing-instances instance-name provider-tunnel selective]
user@host# set group multicast IP address source source IP address
inter-region-segmented fan-out fan-out value
user@host# set group multicast IP address source source IP address
inter-region-segmented threshold rate-value
```
    - Specify fan-out and threshold values for a wildcard source.
 

```
[edit routing-instances instance-name provider-tunnel selective]
user@host# set group multicast IP address source wildcard-source
inter-region-segmented fan-out fan-out value
user@host# set group multicast IP address source wildcard-source
inter-region-segmented threshold rate-value
```
  - Configure inter-region-segmented fan-out value for wildcard-group-inet belonging to group.
 

```
[edit routing-instances instance-name provider-tunnel selective]
user@host# set wildcard-group-inet wildcard-source inter-region-segmented
fan-out fan-out value
```
  - Configure inter-region-segmented fan-out value for wildcard-group-inet6 belonging to group.
 

```
[edit routing-instances instance-name provider-tunnel selective]
user@host# set wildcard-group-inet6 wildcard-source inter-region-segmented
fan-out fan-out value
```
2. Configure the inter-region template on the transit ABR to specify the tunnel type to be used for a specific region or for all regions.
  - Configure the inter-region template to specify the tunnel type such as ingress-replication, ldp-p2mp, and rsvp-te for a specific region.
  - Specify create-new-ucast-tunnel or label-switched-path for tunnel type ingress-replication for a specific region.
 

```
[edit protocols mvpn]
user@host# set inter-region-template template template-name region region-name
ingress-replication create-new-ucast-tunnel
```



```
user@host# set inter-region-template template template-name region region-name
ingress-replication label-switched-path label-switched-path-template (
default-template | default-template)
```

- Specify tunnel type ldp-p2mp for a specific region.

```
[edit protocols mvpn]
user@host# set inter-region-template template template-name region region-name
ldp-p2mp
```

- Specify static lsp or template for label-switched-path-template for tunnel type rsvp-te belonging to a specific region.

```
[edit protocols mvpn]
user@host# set inter-region-template template template-name region region-name
rsvp-te label-switched-path-template (default | lsp-template-name)
user@host# set inter-region-template template template-name region region-name
rsvp-te static-lsp static-lsp
```

- Configure the inter-region template to specify the tunnel type such as ingress-replication, ldp-p2mp, and rsvp-te for all regions.
- Specify create-new-ucast-tunnel or label-switched-path for tunnel type ingress-replication for all regions.

```
[edit protocols mvpn]
user@host# set inter-region-template template template-name all-regions
region-name ingress-replication create-new-ucast-tunnel
user@host# set inter-region-template template template-name all-regions
region-name ingress-replication label-switched-path
label-switched-path-template ( default-template | default-template)
```

- Specify tunnel type ldp-p2mp for all regions.

```
[edit protocols mvpn]
user@host# set inter-region-template template template-name all-regions
region-name ldp-p2mp
```

- Specify static lsp or template for label-switched-path-template for tunnel type rsvp-te belonging to all regions.

```
[edit protocols mvpn]
user@host# set inter-region-template template template-name all-regions
region-name rsvp-te label-switched-path-template (default |
lsp-template-name)
user@host# set inter-region-template template template-name all-regions
region-name rsvp-te static-lsp static-lsp
```

3. Specify the template, which indicates the tunnel types, to be used in inter-region segmentation on the Transit ABRs.

```
[edit routing-instances instance-name provider-tunnel]
user@host# set inter-region template template-name
```

4. Specify no inter-region segmentation if you do not want the ABR to participate in the inter-region segmentation.

```
[edit routing-instances instance-name provider-tunnel]
user@host# set inter-region no-inter-region-segmentation
```

**Related Documentation**

- [Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471](#)
- [Example: Configuring Segmented Inter-Area P2MP LSP on page 478](#)
- [all-regions on page 1042](#)
- [inter-region on page 1082](#)
- [inter-region-template on page 1085](#)
- [inter-region-segmented on page 1083](#)
- [region on page 1122](#)
- [template on page 1141](#)

---

## Example: Configuring Segmented Inter-Area P2MP LSP

This example shows how to segment the P2MP LSPs at the area boundary as described in Internet draft *draft-ietf-mpls-seamless-mcast-14.txt*. You can configure policies on the segmented next-hop extended community (S-NH EC) so that S-PMSI A-D routes with the S-NH EC is reflected by the ABR while all other routes are reflected by other route reflectors.

- [Requirements on page 478](#)
- [Overview on page 478](#)
- [Configuration on page 480](#)
- [Verification on page 534](#)

### Requirements

This example uses the following hardware and software components:

- Fourteen MX Series 3D Universal Edge Routers
- Junos OS Release 15.1 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure OSPF.

### Overview

Starting with Junos OS Release 15.1, P2MP LSPs can be segmented at the area boundary. A segmented P2MP LSP consists of ingress area segment (Ingress PE router or ASBR), backbone area segment (Transit ABR), and egress area segment (Egress PE routers or ASBRs). Each of the intra-area segments can be carried over provider tunnels such as P2MP RSVP-TE LSP, P2MP mLDP LSP, or ingress replication. Segmentation of inter-area P2MP LSP occurs when the S-PMSI autodiscovery (AD) routes are advertised, which triggers the inclusion of a new BGP extended community or inter-area P2MP segmented

To configure inter-region segmentation at the ingress PE router, configure the **inter-region-segmented** statement at the **[edit routing-instances *instance-name* provider-tunnel]** hierarchy level. To configure the inter-region template at the transit ABRs, configure the **inter-region-template *template-name*** statement at the **[edit protocols **mvpn**]** hierarchy level. To configure inter-region segmentation at the transit ABR, configure the **inter-region** statement at the **[edit routing-instance *instance-name* provider-tunnel]** hierarchy level.

In the topology shown in [Figure 53 on page 479](#), the segmented tunnel combination is as follows:

- Ingress area tunnel — PE1 to ABR1 with IR as the tunnel.
- Backbone area tunnel — ABR1, ABR2, and ABR3 with RSVP-TE as the tunnel.
- Egress area tunnel — ABR2 to PE2 and PE4, ABR3 to PE3 with RSVP-TE as the tunnel.

The diagram illustrates a network topology with the following components:

- Table 1: IP Addressing**
- Table 2: Interface Configurations**

IoO:	
PE1	198.51.100.1/24
CE1	198.51.100.16/24
P1	10.255.162.119/32
ABR1	10.255.162.100/32
ABR2	10.255.162.117/32
P2	198.51.100.19/24
ABR3	10.255.162.107/32

Interface Configurations	
PE3	10.255.162.102/32
CE4	198.51.100.21/24
CE5	198.51.100.22/24
PE2	10.255.162.104/32
CE2	198.51.100.23/24
PE4	198.51.100.17/24
CE3	198.51.100.18/24

## 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.

```

PE1 set interfaces ge-2/0/10 unit 1 family inet address 192.0.2.2/24
    set interfaces ge-2/0/10 unit 1 family inet6 address ::192.0.2.2/120
    set interfaces ge-2/0/10 unit 1 family mpls
    set interfaces ge-3/3/2 unit 0 family inet address 203.0.113.1/24
    set interfaces ge-3/3/2 unit 0 family iso
    set interfaces ge-3/3/2 unit 0 family inet6 address ::203.0.113.1/120
    set interfaces ge-3/3/2 unit 0 family mpls
    set interfaces lo0 unit 201 family inet address 198.51.100.1/24
    set routing-options autonomous-system 65550
    set protocols rsvp interface fxp0.0 disable
    set protocols rsvp interface ge-3/3/2.0
    set protocols rsvp interface lo0.0
    set protocols mpls ipv6-tunneling
    set protocols mpls interface fxp0.0 disable
    set protocols mpls interface ge-3/3/2.0
    set protocols mpls interface lo0.0
    set protocols bgp group IBGP type internal
    set protocols bgp group IBGP local-address 10.255.162.109
    set protocols bgp group IBGP family inet any
    set protocols bgp group IBGP family inet-vpn unicast
    set protocols bgp group IBGP family inet-vpn multicast
    set protocols bgp group IBGP family inet6 any
    set protocols bgp group IBGP family inet6-vpn unicast
    set protocols bgp group IBGP family inet6-mvpn signaling
    set protocols bgp group IBGP family inet6-mvpn signaling
    set protocols bgp group IBGP family inet-mdt signaling
    set protocols bgp group IBGP neighbor 10.255.162.100
    set protocols ospf traffic-engineering
    set protocols ospf area 0.0.0.1 interface fxp0.0 disable
    set protocols ospf area 0.0.0.1 interface ge-3/3/2.0
    set protocols ospf area 0.0.0.1 interface lo0.0
    set protocols ldp interface all
    set protocols ldp p2mp
    set protocols pim interface all
    set protocols pim interface fxp0.0 disable
    set protocols pim interface lo0.0
    set protocols pim default-vpn-source interface-name lo0.0
    set policy-options policy-statement bgp-to-ospf from protocol bgp
    set policy-options policy-statement bgp-to-ospf then accept
    set routing-instances vpn1 instance-type vrf
    set routing-instances vpn1 interface ge-2/0/10
    set routing-instances vpn1 interface lo0.201
    set routing-instances vpn1 route-distinguisher 10.255.162.109:100
    set routing-instances vpn1 provider-tunnel selective group 192.0.2.2/24 source 172.16.1.2/32
        ingress-replication label-switched-path
    set routing-instances vpn1 provider-tunnel selective group 192.0.2.2/24 source 172.16.1.2/32
        threshold-rate 10

```

```

set routing-instances vpn1 provider-tunnel selective group 192.0.2.2/24 source 172.16.1.2/32
inter-region-segmented threshold 0
set routing-instances vpn1 provider-tunnel selective group 192.0.2.1/24 source 172.16.1.2/32
ingress-replication label-switched-path
set routing-instances vpn1 provider-tunnel selective group 192.0.2.1/24 source 172.16.1.2/32
threshold-rate 0
set routing-instances vpn1 provider-tunnel selective group 192.0.2.1/24 source 172.16.1.2/32
inter-region-segmented threshold 10
set routing-instances vpn1 provider-tunnel selective group 192.0.2.3/24 source 172.16.1.2/32
ingress-replication label-switched-path
set routing-instances vpn1 provider-tunnel selective group 192.0.2.3/24 source 172.16.1.2/32
threshold-rate 0
set routing-instances vpn1 provider-tunnel selective group 192.0.2.3/24 source 172.16.1.2/32
inter-region-segmented threshold 0
set routing-instances vpn1 provider-tunnel family inet ingress-replication
label-switched-path
set routing-instances vpn1 provider-tunnel family inet6 ingress-replication
label-switched-path
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 protocols ospf export bgp-to-ospf
set routing-instances vpn1 protocols ospf area 0.0.0.1 interface all
set routing-instances vpn1 protocols ospf area 0.0.0.1 interface lo0.201
set routing-instances vpn1 protocols ospf3 export bgp-to-ospf
set routing-instances vpn1 protocols ospf3 area 0.0.0.1 interface all
set routing-instances vpn1 protocols pim dense-groups 192.0.2.39/24
set routing-instances vpn1 protocols pim dense-groups 192.0.2.40/24
set routing-instances vpn1 protocols pim rp local family inet address 198.51.100.1
set routing-instances vpn1 protocols pim rp static address ::198.51.100.1
set routing-instances vpn1 protocols pim interface all mode sparse-dense

```

```

CE1  set interfaces ge-2/0/3 unit 0 family inet address 172.16.1.1/24
set interfaces ge-2/0/3 unit 0 family iso
set interfaces ge-2/0/3 unit 0 family inet6 address 0000:0000:0000:0000:172:2:1:1/120
set interfaces ge-2/0/3 unit 0 family mpls
set interfaces ge-2/0/10 unit 101 family inet address 192.0.2.1/24
set interfaces ge-2/0/10 unit 101 family inet6 address ::192.0.2.1/120
set interfaces lo0 unit 1 family inet address 198.51.100.16/24
set interfaces lo0 unit 1 family inet6 address abcd::198:51:100:16/128
set protocols igmp interface ge-2/0/3.0 version 3
set protocols ospf area 0.0.0.1 interface all
set protocols ospf3 area 0.0.0.1 interface all
set protocols pim dense-groups 192.0.2.39/24
set protocols pim dense-groups 192.0.2.40/24
set protocols pim rp static address 198.51.100.1
set protocols pim rp static address ::198.51.100.1
set protocols pim interface all mode sparse-dense
set protocols pim interface ge-2/0/10.101
set protocols pim interface ge-2/0/3.0
set protocols pim interface lo0.1

P1  set interfaces ge-1/0/5 unit 0 family inet address 203.0.113.2/24
set interfaces ge-1/0/5 unit 0 family iso
set interfaces ge-1/0/5 unit 0 family inet6 address ::203.0.113.2/120

```

```

set interfaces ge-1/0/5 unit 0 family mpls
set interfaces ge-1/0/6 unit 0 family inet address 203.0.113.3/24
set interfaces ge-1/0/6 unit 0 family iso
set interfaces ge-1/0/6 unit 0 family inet6 address ::203.0.113.3/120
set interfaces ge-1/0/6 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 203.0.113.0/24
set interfaces lo0 unit 0 family inet address 10.255.162.119/32 primary
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface ge-1/0/5.0
set protocols rsvp interface ge-1/0/6.0
set protocols rsvp interface lo0.0
set protocols mpls ipv6-tunneling
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-1/0/5.0
set protocols mpls interface ge-1/0/6.0
set protocols mpls interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.1 interface fxp0.0 disable
set protocols ospf area 0.0.0.1 interface ge-1/0/5.0
set protocols ospf area 0.0.0.1 interface ge-1/0/6.0
set protocols ospf area 0.0.0.1 interface lo0.0
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface lo0.0

```

```

ABR1 set interfaces ge-1/0/4 unit 0 family inet address 203.0.113.12/24
set interfaces ge-1/0/4 unit 0 family iso
set interfaces ge-1/0/4 unit 0 family inet6 address ::203.0.113.12/120
set interfaces ge-1/0/4 unit 0 family mpls
set interfaces ge-1/1/3 unit 0 family inet address 203.0.113.5/24
set interfaces ge-1/1/3 unit 0 family iso
set interfaces ge-1/1/3 unit 0 family inet6 address ::203.0.113.5/120
set interfaces ge-1/1/9 unit 0 family inet address 203.0.113.4/24
set interfaces ge-1/1/9 unit 0 family iso
set interfaces ge-1/1/9 unit 0 family inet6 address ::203.0.113.4/120
set interfaces ge-1/1/9 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 203.0.113.0/24
set interfaces lo0 unit 0 family inet address 10.255.162.100/32 primary
set routing-options autonomous-system 65550
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface ge-1/1/9.0
set protocols rsvp interface ge-1/0/4.0
set protocols rsvp interface ge-1/1/3.0
set protocols rsvp interface lo0.0
set protocols rsvp interface all
set protocols mpls ipv6-tunneling
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-1/1/9.0
set protocols mpls interface ge-1/0/4.0
set protocols mpls interface ge-1/1/3.0
set protocols mpls interface lo0.0
set protocols mpls interface all

```

```

set protocols bgp group IBGP_1 type internal
set protocols bgp group IBGP_1 local-address 10.255.162.100
set protocols bgp group IBGP_1 family inet any
set protocols bgp group IBGP_1 family inet-vpn unicast
set protocols bgp group IBGP_1 family inet-vpn multicast
set protocols bgp group IBGP_1 family inet6 any
set protocols bgp group IBGP_1 family inet6-vpn unicast
set protocols bgp group IBGP_1 family inet6-mvpn signaling
set protocols bgp group IBGP_1 family inet6-mvpn signaling
set protocols bgp group IBGP_1 family inet-mdt signaling
set protocols bgp group IBGP_1 cluster 0.0.0.1
set protocols bgp group IBGP_1 neighbor 10.255.162.109
set protocols bgp group IBGP_0 type internal
set protocols bgp group IBGP_0 local-address 10.255.162.100
set protocols bgp group IBGP_0 family inet any
set protocols bgp group IBGP_0 family inet-vpn unicast
set protocols bgp group IBGP_0 family inet-vpn multicast
set protocols bgp group IBGP_0 family inet6 any
set protocols bgp group IBGP_0 family inet6-vpn unicast
set protocols bgp group IBGP_0 family inet6-mvpn signaling
set protocols bgp group IBGP_0 family inet6-mvpn signaling
set protocols bgp group IBGP_0 family inet-mdt signaling
set protocols bgp group IBGP_0 neighbor 10.255.162.117
set protocols bgp group IBGP_0 neighbor 10.255.162.107
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.1 interface fxp0.0 disable
set protocols ospf area 0.0.0.1 interface ge-1/1/9.0
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0
set protocols ospf area 0.0.0.0 interface ge-1/1/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface lo0.0
set protocols mvpn inter-region-template template template_1 region IBGP_0 rsvp-te
  label-switched-path-template default-template
set protocols mvpn inter-region-template template template_2 region IBGP_0 ldp-p2mp
set protocols mvpn inter-region-template template template_3 region IBGP_0
  ingress-replication create-new-ucast-tunnel
set protocols mvpn inter-region-template template template_3 region IBGP_0
  ingress-replication label-switched-path label-switched-path-template default-template
set protocols mvpn inter-region-template template template_4 all-regions incoming
set protocols mvpn inter-region-template template template_5 region IBGP_0 rsvp-te
  static-lsp ABR1_to_ABR3
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 route-distinguisher 10.255.162.100:100
set routing-instances vpn1 provider-tunnel inter-region template template_1
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label

ABR2
set interfaces ge-1/0/3 unit 0 family inet address 203.0.113.11/24
set interfaces ge-1/0/3 unit 0 family iso
set interfaces ge-1/0/3 unit 0 family inet6 address ::203.0.113.11/120
set interfaces ge-1/0/3 unit 0 family mpls

```

```
set interfaces ge-1/1/4 unit 0 family inet address 203.0.113.10/24
set interfaces ge-1/1/4 unit 0 family iso
set interfaces ge-1/1/4 unit 0 family inet6 address ::203.0.113.10/120
set interfaces ge-1/1/4 unit 0 family mpls
set interfaces ge-1/1/10 unit 1 family inet address 192.0.2.2/24
set interfaces ge-1/1/10 unit 1 family inet6 address ::192.0.2.2/120
set interfaces ge-1/1/10 unit 1 family mpls
set interfaces lo0 unit 0 family inet address 203.0.113.0/24
set interfaces lo0 unit 0 family inet address 10.255.162.117/32 primary
set routing-options autonomous-system 65550
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface lo0.0
set protocols rsvp interface all
set protocols mpls ipv6-tunneling
set protocols mpls interface fxp0.0 disable
set protocols mpls interface lo0.0
set protocols mpls interface all
set protocols bgp group IBGP_2 type internal
set protocols bgp group IBGP_2 local-address 10.255.162.117
set protocols bgp group IBGP_2 family inet any
set protocols bgp group IBGP_2 family inet-vpn unicast
set protocols bgp group IBGP_2 family inet-vpn multicast
set protocols bgp group IBGP_2 family inet6 any
set protocols bgp group IBGP_2 family inet6-vpn unicast
set protocols bgp group IBGP_2 family inet6-mvpn signaling
set protocols bgp group IBGP_2 family inet6-mvpn signaling
set protocols bgp group IBGP_2 family inet-mdt signaling
set protocols bgp group IBGP_2 cluster 0.0.0.2
set protocols bgp group IBGP_2 neighbor 10.255.162.104
set protocols bgp group IBGP_2 neighbor 198.51.100.17
set protocols bgp group IBGP_0 type internal
set protocols bgp group IBGP_0 local-address 10.255.162.117
set protocols bgp group IBGP_0 family inet any
set protocols bgp group IBGP_0 family inet-vpn unicast
set protocols bgp group IBGP_0 family inet-vpn multicast
set protocols bgp group IBGP_0 family inet6 any
set protocols bgp group IBGP_0 family inet6-vpn unicast
set protocols bgp group IBGP_0 family inet6-mvpn signaling
set protocols bgp group IBGP_0 family inet6-mvpn signaling
set protocols bgp group IBGP_0 family inet-mdt signaling
set protocols bgp group IBGP_0 neighbor 10.255.162.100
set protocols bgp group IBGP_0 neighbor 10.255.162.107
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-1/0/3.0
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.2 interface ge-1/1/10.1
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface lo0.0
set protocols mvpn inter-region-template template template_1 region IBGP_2 rsvp-te
  label-switched-path-template default-template
set protocols mvpn inter-region-template template template_2 region IBGP_2 ldp-p2mp
```



```

set protocols mvpn inter-region-template template template_3 region IBGP_2
  ingress-replication create-new-ucast-tunnel
set protocols mvpn inter-region-template template template_3 region IBGP_2
  ingress-replication label-switched-path label-switched-path-template default-template
set protocols mvpn inter-region-template template template_4 all-regions incoming
set protocols mvpn inter-region-template template template_5 region IBGP_2 rsvp-te
  static-lsp ABR2_to_PE2_3
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 route-distinguisher 10.255.162.117:100
set routing-instances vpn1 provider-tunnel inter-region template template_1
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label

P2    set interfaces ge-1/0/5 unit 0 family inet address 203.0.113.14/24
      set interfaces ge-1/0/5 unit 0 family iso
      set interfaces ge-1/0/5 unit 0 family inet6 address ::203.0.113.14/120
      set interfaces ge-1/0/5 unit 0 family mpls
      set interfaces ge-1/1/10 unit 101 family inet address 192.0.2.1/24
      set interfaces ge-1/1/10 unit 101 family inet6 address ::192.0.2.1/120
      set interfaces ge-1/1/10 unit 101 family mpls
      set interfaces lo0 unit 1 family inet address 198.51.100.19/24
      set interfaces lo0 unit 1 family inet6 address abcd::198:51:100:19/128
      set protocols rsvp interface fxp0.0 disable
      set protocols rsvp interface all
      set protocols rsvp interface lo0.1
      set protocols mpls ipv6-tunneling
      set protocols mpls interface fxp0.0 disable
      set protocols mpls interface all
      set protocols mpls interface lo0.1
      set protocols ospf traffic-engineering
      set protocols ospf area 0.0.0.2 interface all
      set protocols ospf area 0.0.0.2 interface lo0.1
      set protocols ldp interface all
      set protocols ldp p2mp
      set protocols pim dense-groups 192.0.2.39/24
      set protocols pim dense-groups 192.0.2.40/24
      set protocols pim rp static address 198.51.100.1
      set protocols pim rp static address ::198.51.100.1
      set protocols pim interface all mode sparse-dense

ABR3  set interfaces ge-1/0/8 unit 0 family inet address 203.0.113.6/24
      set interfaces ge-1/0/8 unit 0 family iso
      set interfaces ge-1/0/8 unit 0 family inet6 address ::203.0.113.6/120
      set interfaces ge-1/0/8 unit 0 family mpls
      set interfaces ge-1/1/4 unit 0 family inet address 203.0.113.9/24
      set interfaces ge-1/1/4 unit 0 family iso
      set interfaces ge-1/1/4 unit 0 family inet6 address ::203.0.113.9/120
      set interfaces ge-1/1/4 unit 0 family mpls
      set interfaces ge-1/3/1 unit 0 family inet address 203.0.113.8/24
      set interfaces ge-1/3/1 unit 0 family iso
      set interfaces ge-1/3/1 unit 0 family inet6 address ::203.0.113.8/120
      set interfaces ge-1/3/1 unit 0 family mpls
      set interfaces lo0 unit 0 family inet address 203.0.113.0/24
      set interfaces lo0 unit 0 family inet address 10.255.162.107/32 primary

```

```
set routing-options autonomous-system 65550
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface lo0.0
set protocols mpls ipv6-tunneling
set protocols mpls label-switched-path ABR3_to_PE3 from 10.255.162.107
set protocols mpls label-switched-path ABR3_to_PE3 to 10.255.162.102
set protocols mpls label-switched-path ABR3_to_PE3 p2mp vpn1
set protocols mpls label-switched-path ABR3_to_ABR1 from 10.255.162.107
set protocols mpls label-switched-path ABR3_to_ABR1 to 10.255.162.100
set protocols mpls label-switched-path ABR3_to_ABR1 p2mp vpn1
set protocols mpls label-switched-path ABR3_to_ABR2 from 10.255.162.107
set protocols mpls label-switched-path ABR3_to_ABR2 to 10.255.162.117
set protocols mpls label-switched-path ABR3_to_ABR2 p2mp vpn1
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface lo0.0
set protocols bgp group IBGP_3 type internal
set protocols bgp group IBGP_3 local-address 10.255.162.107
set protocols bgp group IBGP_3 family inet any
set protocols bgp group IBGP_3 family inet-vpn unicast
set protocols bgp group IBGP_3 family inet-vpn multicast
set protocols bgp group IBGP_3 family inet6 any
set protocols bgp group IBGP_3 family inet6-vpn unicast
set protocols bgp group IBGP_3 family inet6-mvpn signaling
set protocols bgp group IBGP_3 family inet6-mvpn signaling
set protocols bgp group IBGP_3 family inet-mdt signaling
set protocols bgp group IBGP_3 cluster 0.0.0.3
set protocols bgp group IBGP_3 neighbor 10.255.162.102
set protocols bgp group IBGP_0 type internal
set protocols bgp group IBGP_0 local-address 10.255.162.107
set protocols bgp group IBGP_0 family inet any
set protocols bgp group IBGP_0 family inet-vpn unicast
set protocols bgp group IBGP_0 family inet-vpn multicast
set protocols bgp group IBGP_0 family inet6 any
set protocols bgp group IBGP_0 family inet6-vpn unicast
set protocols bgp group IBGP_0 family inet6-mvpn signaling
set protocols bgp group IBGP_0 family inet6-mvpn signaling
set protocols bgp group IBGP_0 family inet-mdt signaling
set protocols bgp group IBGP_0 neighbor 10.255.162.100
set protocols bgp group IBGP_0 neighbor 10.255.162.117
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-1/0/8.0
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.3 interface ge-1/3/1.0
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface lo0.0
set protocols mvpn inter-region-template template template_1 region IBGP_3 rsvp-te
  label-switched-path-template default-template
set protocols mvpn inter-region-template template template_2 region IBGP_3 ldp-p2mp
```

```

set protocols mvpn inter-region-template template template_3 region IBGP_3
  ingress-replication create-new-ucast-tunnel
set protocols mvpn inter-region-template template template_3 region IBGP_3
  ingress-replication label-switched-path label-switched-path-template default-template
set protocols mvpn inter-region-template template template_4 all-regions incoming
set protocols mvpn inter-region-template template template_5 region IBGP_3 rsvp-te
  static-lsp ABR3_to_PE3
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 route-distinguisher 10.255.162.107:100
set routing-instances vpn1 provider-tunnel inter-region template template_1
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label

```

```

PE3  set interfaces ge-3/0/1 unit 0 family inet address 203.0.113.15/24
      set interfaces ge-3/0/1 unit 0 family iso
      set interfaces ge-3/0/1 unit 0 family inet6 address ::203.0.113.15/120
      set interfaces ge-3/0/1 unit 0 family mpls
      set interfaces ge-3/1/0 unit 0 family inet address 203.0.113.7/24
      set interfaces ge-3/1/0 unit 0 family iso
      set interfaces ge-3/1/0 unit 0 family inet6 address ::203.0.113.7/120
      set interfaces ge-3/1/0 unit 0 family mpls
      set interfaces ge-3/2/0 unit 1 family inet address 192.0.2.9/24
      set interfaces ge-3/2/0 unit 1 family inet6 address ::192.0.2.9/120
      set interfaces ge-3/2/0 unit 1 family mpls
      set interfaces ge-3/2/0 unit 2 family inet address 192.0.2.10/24
      set interfaces ge-3/2/0 unit 2 family inet6 address ::192.0.2.10/120
      set interfaces ge-3/2/0 unit 2 family mpls
      set interfaces lo0 unit 0 family inet address 203.0.113.0/24
      set interfaces lo0 unit 0 family inet address 10.255.162.102/32 primary
      set routing-options autonomous-system 65550
      set protocols rsvp interface fxp0.0 disable
      set protocols rsvp interface all
      set protocols rsvp interface lo0.0
      set protocols mpls ipv6-tunneling
      set protocols mpls label-switched-path PE3_to_PE2 from 10.255.162.102
      set protocols mpls label-switched-path PE3_to_PE2 to 10.255.162.104
      set protocols mpls label-switched-path PE3_to_PE2 p2mp vpn1
      set protocols mpls label-switched-path PE3_to_PE4 from 10.255.162.102
      set protocols mpls label-switched-path PE3_to_PE4 to 198.51.100.17
      set protocols mpls label-switched-path PE3_to_PE4 p2mp vpn1
      set protocols mpls label-switched-path PE3_to_PE1 from 10.255.162.102
      set protocols mpls label-switched-path PE3_to_PE1 to 10.255.162.109
      set protocols mpls label-switched-path PE3_to_PE1 p2mp vpn1
      set protocols mpls label-switched-path PE3_to_ABR3 from 10.255.162.102
      set protocols mpls label-switched-path PE3_to_ABR3 to 10.255.162.107
      set protocols mpls label-switched-path PE3_to_ABR3 p2mp vpn1
      set protocols mpls interface fxp0.0 disable
      set protocols mpls interface all
      set protocols mpls interface lo0.0
      set protocols bgp group IBGP type internal
      set protocols bgp group IBGP local-address 10.255.162.102
      set protocols bgp group IBGP family inet any
      set protocols bgp group IBGP family inet-vpn unicast
      set protocols bgp group IBGP family inet-vpn multicast
      set protocols bgp group IBGP family inet6 any

```

```

set protocols bgp group IBGP family inet6-vpn unicast
set protocols bgp group IBGP family inet-mvpn signaling
set protocols bgp group IBGP family inet6-mvpn signaling
set protocols bgp group IBGP family inet-mdt signaling
set protocols bgp group IBGP neighbor 10.255.162.107
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.3 interface fxp0.0 disable
set protocols ospf area 0.0.0.3 interface all
set protocols ospf area 0.0.0.3 interface lo0.0
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface lo0.0
set protocols pim default-vpn-source interface-name lo0.0
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-3/2/0.1
set routing-instances vpn1 interface ge-3/2/0.2
set routing-instances vpn1 route-distinguisher 10.255.162.102:100
set routing-instances vpn1 provider-tunnel family inet ingress-replication
    label-switched-path
set routing-instances vpn1 provider-tunnel family inet6 ingress-replication
    label-switched-path
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 protocols ospf export bgp-to-ospf
set routing-instances vpn1 protocols ospf area 0.0.0.3 interface all
set routing-instances vpn1 protocols ospf3 export bgp-to-ospf
set routing-instances vpn1 protocols ospf3 area 0.0.0.3 interface all
set routing-instances vpn1 protocols pim dense-groups 192.0.2.39/24
set routing-instances vpn1 protocols pim dense-groups 192.0.2.40/24
set routing-instances vpn1 protocols pim rp static address 198.51.100.1
set routing-instances vpn1 protocols pim rp static address ::198.51.100.1
set routing-instances vpn1 protocols pim interface all mode sparse-dense
set routing-instances vpn1 protocols mvpn mvpn-mode spt-only

```

```

CE4 set interfaces ge-3/1/1 unit 0 family inet address 172.16.0.1/24
set interfaces ge-3/1/1 unit 0 family iso
set interfaces ge-3/1/1 unit 0 family inet6 address 0000:0000:0000:0000:172:16:0:1/120
set interfaces ge-3/1/1 unit 0 family mpls
set interfaces ge-3/2/0 unit 102 description "Link to PE3_1 from CE3_2"
set interfaces ge-3/2/0 unit 102 family inet address 192.0.2.11/24
set interfaces ge-3/2/0 unit 102 family inet6 address ::192.0.2.11/120
set interfaces ge-3/2/0 unit 102 family mpls
set interfaces lo0 unit 2 family inet address 198.51.100.21/24
set interfaces lo0 unit 2 family inet6 address abcd::198:51:100:21/128
set protocols igmp interface ge-3/1/1.0 version 3
set protocols mld interface ge-3/1/1.0 version 2
set protocols ospf area 0.0.0.3 interface all
set protocols ospf3 area 0.0.0.3 interface all
set protocols pim dense-groups 192.0.2.39/24
set protocols pim dense-groups 192.0.2.40/24
set protocols pim rp static address 198.51.100.1

```

```

set protocols pim rp static address ::198.51.100.1
set protocols pim interface all mode sparse-dense

CE5
set interfaces ge-3/2/0 unit 101 family inet address 192.0.2.8/24
set interfaces ge-3/2/0 unit 101 family inet6 address ::192.0.2.8/120
set interfaces ge-3/2/0 unit 101 family mpls
set interfaces lo0 unit 1 family inet address 198.51.100.22/24
set interfaces lo0 unit 1 family inet6 address abcd::198:51:100:22/128
set protocols ospf area 0.0.0.3 interface all
set protocols ospf3 area 0.0.0.3 interface all
set protocols pim dense-groups 192.0.2.39/24
set protocols pim dense-groups 192.0.2.40/24
set protocols pim rp static address 198.51.100.1
set protocols pim rp static address ::198.51.100.1
set protocols pim interface all mode sparse-dense

PE2
set interfaces ge-1/0/5 unit 0 family inet address 203.0.113.13/24
set interfaces ge-1/0/5 unit 0 family iso
set interfaces ge-1/0/5 unit 0 family inet6 address ::203.0.113.13/120
set interfaces ge-1/0/5 unit 0 family mpls
set interfaces ge-1/1/00 unit 1 family inet address 192.0.2.4/24
set interfaces ge-1/1/0 unit 1 family inet6 address ::192.0.2.4/120
set interfaces ge-1/1/0 unit 1 family mpls
set interfaces ge-1/1/0 unit 2 family inet address 192.0.2.12/24
set interfaces ge-1/1/0 unit 2 family inet6 address ::192.0.2.12/120
set interfaces ge-1/1/0 unit 2 family mpls
set interfaces vt-1/1/0 unit 1 family inet
set interfaces vt-1/1/0 unit 1 family inet6
set interfaces lo0 unit 0 family inet address 203.0.113.0/24
set interfaces lo0 unit 0 family inet address 10.255.162.104/24 primary
set interfaces lo0 unit 201 family inet6 address ::198.51.100.1/128
set routing-options autonomous-system 65550
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface all
set protocols rsvp interface lo0.0
set protocols mpls ipv6-tunneling
set protocols mpls label-switched-path PE2_to_PE3 from 10.255.162.104
set protocols mpls label-switched-path PE2_to_PE3 to 10.255.162.102
set protocols mpls label-switched-path PE2_to_PE3 p2mp vpn1
set protocols mpls label-switched-path PE2_to_PE4 from 10.255.162.104
set protocols mpls label-switched-path PE2_to_PE4 to 198.51.100.17
set protocols mpls label-switched-path PE2_to_PE4 p2mp vpn1
set protocols mpls label-switched-path PE2_to_PE1 from 10.255.162.104
set protocols mpls label-switched-path PE2_to_PE1 to 10.255.162.109
set protocols mpls label-switched-path PE2_to_PE1 p2mp vpn1
set protocols mpls label-switched-path PE2_to_ABR2 from 10.255.162.104
set protocols mpls label-switched-path PE2_to_ABR2 to 10.255.162.117
set protocols mpls label-switched-path PE2_to_ABR2 p2mp vpn1
set protocols mpls interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface lo0.0
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 10.255.162.104
set protocols bgp group IBGP family inet any

```

```

set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP family inet-vpn multicast
set protocols bgp group IBGP family inet6 any
set protocols bgp group IBGP family inet6-vpn unicast
set protocols bgp group IBGP family inet-mvpn signaling
set protocols bgp group IBGP family inet6-mvpn signaling
set protocols bgp group IBGP family inet-mdt signaling
set protocols bgp group IBGP neighbor 10.255.162.117
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.2 interface fxp0.0 disable
set protocols ospf area 0.0.0.2 interface all
set protocols ospf area 0.0.0.2 interface lo0.0
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface lo0.0
set protocols pim default-vpn-source interface-name lo0.0
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-1/1/0.1
set routing-instances vpn1 interface vt-1/1/0.1 multicast
set routing-instances vpn1 interface lo0.201
set routing-instances vpn1 route-distinguisher 10.255.162.104:100
set routing-instances vpn1 provider-tunnel family inet ingress-replication
    label-switched-path
set routing-instances vpn1 provider-tunnel family inet6 ingress-replication
    label-switched-path
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 protocols ospf export bgp-to-ospf
set routing-instances vpn1 protocols ospf area 0.0.0.2 interface all
set routing-instances vpn1 protocols ospf area 0.0.0.2 interface lo0.201
set routing-instances vpn1 protocols ospf3 export bgp-to-ospf
set routing-instances vpn1 protocols ospf3 area 0.0.0.2 interface all
set routing-instances vpn1 protocols pim dense-groups 192.0.2.39/24
set routing-instances vpn1 protocols pim dense-groups 192.0.2.40/24
set routing-instances vpn1 protocols pim rp local family inet6 address ::198.51.100.1
set routing-instances vpn1 protocols pim rp static address 198.51.100.1
set routing-instances vpn1 protocols pim interface all mode sparse-dense
set routing-instances vpn1 protocols mvpn mvpn-mode spt-only

```

```

CE2 set interfaces ge-1/0/0 unit 0 family inet address 172.17.1.1/24
set l interfaces ge-1/0/0 unit 0 family iso
set interfaces ge-1/0/0 unit 0 family inet6 address 0000:0000:0000:0000:172:17:1:1/120
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/1/0 unit 101 family inet address 192.0.2.3/24
set interfaces ge-1/1/0 unit 101 family inet6 address ::192.0.2.3/120
set interfaces ge-1/1/0 unit 101 family mpls
set interfaces lo0 unit 1 family inet address 198.51.100.23/24
set interfaces lo0 unit 1 family inet6 address abcd::198:51:100:23/128
set protocols igmp interface ge-1/0/0.0 version 3
set protocols mld interface ge-1/0/0.0 version 2
set protocols ospf area 0.0.0.2 interface all

```

```

set protocols ospf3 area 0.0.0.2 interface all
set protocols pim dense-groups 192.0.2.39/24
set protocols pim dense-groups 192.0.2.40/24
set protocols pim rp static address 198.51.100.1
set protocols pim rp static address ::198.51.100.1
set protocols pim interface all mode sparse-dense

```

```

PE4  set interfaces ge-1/1/0 unit 3 family inet address 192.0.2.7/24
      set interfaces ge-1/1/0 unit 3 family inet6 address ::192.0.2.7/120
      set interfaces ge-1/1/0 unit 3 family mpls
      set interfaces ge-1/1/0 unit 102 family inet address 192.0.2.5/24
      set interfaces ge-1/1/0 unit 102 family inet6 address ::192.0.2.5/120
      set interfaces ge-1/1/0 unit 102 family mpls
      set interfaces vt-1/1/0 unit 0 family inet
      set interfaces vt-1/1/0 unit 0 family inet6
      set interfaces lo0 unit 2 family inet address 198.51.100.17/24
      set interfaces lo0 unit 2 family inet6 address abcd::198:51:100:17/128
      set protocols rsvp interface fxp0.0 disable
      set protocols rsvp interface all
      set protocols rsvp interface lo0.2
      set protocols mpls ipv6-tunneling
      set protocols mpls label-switched-path PE4_to_PE3 from 198.51.100.17
      set protocols mpls label-switched-path PE4_to_PE3 to 10.255.162.102
      set protocols mpls label-switched-path PE4_to_PE2 from 198.51.100.17
      set protocols mpls label-switched-path PE4_to_PE2 to 10.255.162.104
      set protocols mpls label-switched-path PE4_to_PE1 from 198.51.100.17
      set protocols mpls label-switched-path PE4_to_PE1 to 10.255.162.109
      set protocols mpls label-switched-path PE4_to_ABR2 from 198.51.100.17
      set protocols mpls label-switched-path PE4_to_ABR2 to 10.255.162.117
      set protocols mpls interface fxp0.0 disable
      set protocols mpls interface all
      set protocols mpls interface lo0.2
      set protocols bgp group IBGP type internal
      set protocols bgp group IBGP local-address 198.51.100.17
      set protocols bgp group IBGP family inet any
      set protocols bgp group IBGP family inet-vpn unicast
      set protocols bgp group IBGP family inet-vpn multicast
      set protocols bgp group IBGP family inet6 any
      set protocols bgp group IBGP family inet6-vpn unicast
      set protocols bgp group IBGP family inet6-mvpn signaling
      set protocols bgp group IBGP family inet6-mvpn signaling
      set protocols bgp group IBGP family inet-mdt signaling
      set protocols bgp group IBGP neighbor 10.255.162.117
      set protocols ospf traffic-engineering
      set protocols ospf area 0.0.0.2 interface fxp0.0 disable
      set protocols ospf area 0.0.0.2 interface all
      set protocols ospf area 0.0.0.2 interface lo0.2
      set protocols ldp interface all
      set protocols ldp p2mp
      set protocols pim interface all
      set protocols pim interface fxp0.0 disable
      set protocols pim interface lo0.2
      set protocols pim default-vpn-source interface-name lo0.2
      set policy-options policy-statement bgp-to-ospf from protocol bgp
      set policy-options policy-statement bgp-to-ospf then accept

```

```

set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface vt-1/1/0.0 multicast
set routing-instances vpn1 interface ge-1/1/0.3
set routing-instances vpn1 route-distinguisher 198.51.100.17:100
set routing-instances vpn1 provider-tunnel family inet ingress-replication
  label-switched-path
set routing-instances vpn1 provider-tunnel family inet6 ingress-replication
  label-switched-path
set routing-instances vpn1 vrf-target target:123:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 protocols ospf export bgp-to-ospf
set routing-instances vpn1 protocols ospf area 0.0.0.2 interface all
set routing-instances vpn1 protocols ospf3 export bgp-to-ospf
set routing-instances vpn1 protocols ospf3 area 0.0.0.2 interface all
set routing-instances vpn1 protocols pim dense-groups 192.0.2.39/24
set routing-instances vpn1 protocols pim dense-groups 192.0.2.40/24
set routing-instances vpn1 protocols pim rp static address 198.51.100.1
set routing-instances vpn1 protocols pim rp static address ::198.51.100.1
set routing-instances vpn1 protocols pim interface all mode sparse-dense
set routing-instances vpn1 protocols mvpn mvpn-mode spt-only
set routing-options autonomous-system 65550

```

```

CE3  set interfaces ge-1/1/0 unit 103 family inet address 192.0.2.6/24
set interfaces ge-1/1/0 unit 103 family inet6 address ::192.0.2.6/120
set interfaces ge-1/1/0 unit 103 family mpls
set interfaces ge-2/1/1 unit 0 family inet address 172.17.2.1/24
set interfaces ge-2/1/1 unit 0 family iso
set interfaces ge-2/1/1 unit 0 family inet6 address 0000:0000:0000:0000:172:17:2:1/120
set interfaces ge-2/1/1 unit 0 family mpls
set interfaces lo0 unit 3 family inet address 198.51.100.18/24
set interfaces lo0 unit 3 family inet6 address abcd::198:51:100:18/128
set protocols igmp interface ge-2/1/1.0 version 3
set protocols mld interface ge-2/1/1.0 version 2
set protocols ospf area 0.0.0.2 interface all
set protocols ospf3 area 0.0.0.2 interface all
set protocols pim dense-groups 192.0.2.39/24
set protocols pim dense-groups 192.0.2.40/24
set protocols pim rp static address 198.51.100.1
set protocols pim rp static address ::198.51.100.1
set protocols pim interface all mode sparse-dense

```

## Configuring PE1

**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 PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set ge-2/0/10 unit 1 family inet address 192.0.2.2/24
user@PE1# set ge-2/0/10 unit 1 family inet6 address ::192.0.2.2/120
user@PE1# set ge-2/0/10 unit 1 family mpls

```



```

user@PE1# set ge-3/3/2 unit 0 family inet address 203.0.113.1/24
user@PE1# set ge-3/3/2 unit 0 family iso
user@PE1# set ge-3/3/2 unit 0 family inet6 address ::203.0.113.1/120
user@PE1# set ge-3/3/2 unit 0 family mpls

```

```

user@PE1# set lo0 unit 201 family inet address 198.51.100.1/24

```

2. Configure the autonomous system number.

```

[edit routing-options]
user@PE1# set autonomous-system 65550

```

3. Disable RSVP on the management interface and enable RSVP on the interfaces.

```

[edit protocols rsvp]
user@PE1# set interface fxp0.0 disable
user@PE1# set interface ge-3/3/2.0
user@PE1# set interface lo0.0

```

4. Enable IPv6 tunneling.

```

[edit protocols mpls]
user@PE1# set ipv6-tunneling

```

5. Disable MPLS on the management interface and enable MPLS on the interfaces.

```

[edit protocols mpls]
user@PE1# set interface fxp0.0 disable
user@PE1# set interface ge-3/3/2.0
user@PE1# set interface lo0.0

```

6. Configure the BGP protocol.

```

[edit protocols bgp]
user@PE1# set group IBGP type internal
user@PE1# set group IBGP local-address 10.255.162.109
user@PE1# set group IBGP family inet any
user@PE1# set group IBGP family inet-vpn unicast
user@PE1# set group IBGP family inet-vpn multicast
user@PE1# set group IBGP family inet6 any
user@PE1# set group IBGP family inet6-vpn unicast
user@PE1# set group IBGP family inet-mvpn signaling
user@PE1# set group IBGP family inet6-mvpn signaling
user@PE1# set group IBGP family inet-mdt signaling
user@PE1# set group IBGP neighbor 10.255.162.100

```

7. Configure OSPF traffic engineering attributes and enable OSPF on the interfaces.

```

[edit protocols ospf]
user@PE1# set traffic-engineering
user@PE1# set area 0.0.0.1 interface fxp0.0 disable
user@PE1# set area 0.0.0.1 interface ge-3/3/2.0
user@PE1# set area 0.0.0.1 interface lo0.0

```

8. Enable LDP on all the interfaces and advertise P2MP capability to peers.

```
[edit protocols ldp]
user@PE1# set interface all
user@PE1# set p2mp
```

9. Configure PIM on the interfaces.

```
[edit protocols pim]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
user@PE1# set interface lo0.0
user@PE1# set default-vpn-source interface-name lo0.0
```

10. Configure the routing policy.

```
[edit policy-options policy-statement]
user@PE1# set bgp-to-ospf from protocol bgp
user@PE1# set bgp-to-ospf then accept
```

11. Configure the routing instance type, interface, and the route distinguisher for the routing instance.

```
[edit routing-instances]
user@PE1# set vpn1 instance-type vrf
user@PE1# set vpn1 interface ge-2/0/10
user@PE1# set vpn1 interface lo0.201
user@PE1# set vpn1 route-distinguisher 10.255.162.109:100
```

12. Configure provider tunnel attributes for the routing instance.

```
[edit routing-instances]
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.2/24 source 172.16.1.2/32
ingress-replication label-switched-path
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.2/24 source 172.16.1.2/32
threshold-rate 10
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.2/24 source 172.16.1.2/32
inter-region-segmented threshold 0

user@PE1# set vpn1 provider-tunnel selective group 192.0.2.1/24 source 172.16.1.2/32
ingress-replication label-switched-path
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.1/24 source 172.16.1.2/32
threshold-rate 0
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.1/24 source 172.16.1.2/32
inter-region-segmented threshold 10

user@PE1# set vpn1 provider-tunnel selective group 192.0.2.3/24 source 172.16.1.2/32
ingress-replication label-switched-path
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.3/24 source 172.16.1.2/32
threshold-rate 0
user@PE1# set vpn1 provider-tunnel selective group 192.0.2.3/24 source 172.16.1.2/32
inter-region-segmented threshold 0
```

```

user@PE1# set vpn1 provider-tunnel family inet ingress-replication
label-switched-path
user@PE1# set vpn1 provider-tunnel family inet6 ingress-replication
label-switched-path

```

13. Configure the VRF target community and advertise a single VPN label for all the routes in the VRF.

```

[edit routing-instances]
user@PE1# set vpn1 vrf-target target:123:1
user@PE1# set vpn1 vrf-table-label

```

14. Enable OSPF for the routing instance.

```

[edit routing-instances]
user@PE1# set vpn1 protocols ospf export bgp-to-ospf
user@PE1# set vpn1 protocols ospf area 0.0.0.1 interface all
user@PE1# set vpn1 protocols ospf area 0.0.0.1 interface lo0.201

```

15. Enable OSPF3 for the routing instance.

```

[edit routing-instances]
user@PE1# set vpn1 protocols ospf3 export bgp-to-ospf
user@PE1# set vpn1 protocols ospf3 area 0.0.0.1 interface all

```

16. Enable PIM attributes for the routing instance.

```

[edit routing-instances]
user@PE1# set vpn1 protocols pim dense-groups 192.0.2.39/24
user@PE1# set vpn1 protocols pim dense-groups 192.0.2.40/24
user@PE1# set vpn1 protocols pim rp local family inet address 198.51.100.1
user@PE1# set vpn1 protocols pim rp static address ::198.51.100.1
user@PE1# set vpn1 protocols pim interface all mode sparse-dense

```

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

```

user@PE1# show interfaces
ge-2/0/10 {
  unit 1 {
    family inet {
      address 192.0.2.2/24;
    }
    family inet6 {
      address ::192.0.2.2/120;
    }
    family mpls;
  }
}
ge-3/3/2 {
  unit 0 {

```

```
family inet {
    address 203.0.113.1/24;
}
family iso;
family inet6 {
    address ::203.0.113.1/120;
}
family mpls;
}
}
lo0 {
    unit 201 {
        family inet {
            address 198.51.100.1/24;
        }
    }
}

user@PE1# show policy-options
policy-statement bgp-to-ospf {
    from protocol bgp;
    then accept;
}

user@PE1# show protocols
rsvp {
    interface fxp0.0 {
        disable;
    }
    interface ge-3/3/2.0;
    interface lo0.0;
}
mpls {
    ipv6-tunneling;
    interface fxp0.0 {
        disable;
    }
    interface ge-3/3/2.0;
    interface lo0.0;
}
bgp {
    group IBGP {
        type internal;
        local-address 10.255.162.109;
        family inet {
            any;
        }
        family inet-vpn {
            unicast;
            multicast;
        }
        family inet6 {
            any;
        }
        family inet6-vpn {
            unicast;
        }
    }
}
```

```

    family inet-mvpn {
        signaling;
    }
    family inet6-mvpn {
        signaling;
    }
    family inet-mdt {
        signaling;
    }
    neighbor 10.255.162.100;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.1 {
        interface fxp0.0 {
            disable;
        }
        interface ge-3/3/2.0;
        interface lo0.0;
    }
}
ldp {
    interface all;
    p2mp;
}
pim {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
    default-vpn-source {
        interface-name lo0.0;
    }
}
}

user@PE1# show routing-instances
vpn1 {
    instance-type vrf;
    interface ge-2/0/10;
    interface lo0.201;
    route-distinguisher 10.255.162.109:100;
    provider-tunnel {
        selective {
            group 192.0.2.2/24 {
                source 172.16.1.2/32 {
                    ingress-replication {
                        label-switched-path;
                    }
                    threshold-rate 10;
                    inter-region-segmented {
                        threshold 0;
                    }
                }
            }
        }
    }
}

```

```
group 192.0.2.1/24 {
  source 172.16.1.2/32 {
    ingress-replication {
      label-switched-path;
    }
    threshold-rate 0;
    inter-region-segmented {
      threshold 10;
    }
  }
}
group 192.0.2.3/24 {
  source 172.16.1.2/32 {
    ingress-replication {
      label-switched-path;
    }
    threshold-rate 0;
    inter-region-segmented {
      threshold 0;
    }
  }
}
family {
  inet {
    ingress-replication {
      label-switched-path;
    }
  }
  inet6 {
    ingress-replication {
      label-switched-path;
    }
  }
}
vrf-target target:123:1;
vrf-table-label;
protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.1 {
      interface all;
      interface lo0.201;
    }
  }
  ospf3 {
    export bgp-to-ospf;
    area 0.0.0.1 {
      interface all;
    }
  }
  pim {
    dense-groups {
      192.0.2.39/24;
      192.0.2.40/24;
```

```

    }
    rp {
      local {
        family inet {
          address 198.51.100.1;
        }
      }
      static {
        address ::198.51.100.1;
      }
    }
    interface all {
      mode sparse-dense;
    }
  }
}

user@PE1# show routing-options
autonomous-system 65550;

```

### Verification

Confirm that the configuration is working properly.

- [Verifying Inflow at the Ingress PE Router on page 499](#)
- [Verifying the Route Table for Segmented Type-3 Traffic Generated from Device ABR1 Toward PE1 Router on page 501](#)
- [Verifying the Route Table for Segmented Type-4 Traffic Received from Device ABR1 Toward PE1 Router on page 503](#)
- [Verifying the LDP Traffic Statistics on page 505](#)

#### *Verifying Inflow at the Ingress PE Router*

**Purpose** Verify the traffic inflow into the ingress PE router for the given routing instance.

**Action** From operational mode, run the **show multicast route extensive instance vpn1** command for Device PE1.

```
user@PE1> show multicast route extensive instance vpn1
display-tunnel-name
Instance: vpn1 Family: INET

Group: 192.0.2.2
  Source: 172.16.1.2/32
  Upstream interface: ge-2/0/10.1
  Downstream interface list:
    mvpn:2
  Number of outgoing interfaces: 1
  Session description: Unknown
  Statistics: 3002 kbps, 10008 pps, 34124622 packets
  Next-hop ID: 0
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Pruned
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:56:53

Group: 192.0.2.1
  Source: 172.16.1.2/32
  Upstream interface: ge-2/0/10.1
  Downstream interface list:
    mvpn:4
  Number of outgoing interfaces: 1
  Session description: Unknown
  Statistics: 3002 kbps, 10008 pps, 34125577 packets
  Next-hop ID: 0
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Pruned
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:56:53

Group: 192.0.2.3
  Source: 172.16.1.2/32
  Upstream interface: ge-2/0/10.1
  Downstream interface list:
    mvpn:3
  Number of outgoing interfaces: 1
  Session description: Unknown
  Statistics: 3002 kbps, 10008 pps, 34124620 packets
  Next-hop ID: 0
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Pruned
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:56:53
```

**Meaning** The output shows the traffic inflow into the ingress Device PE1.



***Verifying the Route Table for Segmented Type-3 Traffic Generated from Device ABR1  
Toward PE1 Router***

**Purpose** Verify the route table for segmented Type-3 traffic generated from Device ABR1.

**Action** From operational mode, run the **show route table vpn1.mvpn.0 match-prefix 3:\* detail** command.

```
user@PE1> show route table vpn1.mvpn.0 match-prefix 3:* detail
```

```
vpn1.mvpn.0: 19 destinations, 22 routes (19 active, 3 holddown, 0 hidden)
3:10.255.162.109:100:32:172.16.1.2:32:20192.0.2.2:10.255.162.109/240 (1 entry, 1 announced)
```

```
    *MVPN    Preference: 70
             PMSI: Flags 0x1: Label 0: Type INGRESS-REPLICATION 10.255.162.109
```

```
             Next hop type: Indirect, Next hop index: 0
             Address: 0xa5b8690
             Next-hop reference count: 11
             Protocol next hop: 10.255.162.109
             Indirect next hop: 0x0 - INH Session ID: 0x0
             State: <Active Int Ext>
             Age: 1:00:20    Metric2: 1
             Validation State: unverified
             Task: mvpn global task
             Announcement bits (3): 0-PIM.vpn1 1-mvpn global task 2-rt-export
```

```
             AS path: I
             Communities: segmented-nh:10.255.162.109:0
```

```
3:10.255.162.109:100:32:172.16.1.2:32:20192.0.2.1:10.255.162.109/240 (1 entry, 1 announced)
```

```
    *MVPN    Preference: 70
             PMSI: Flags 0x1: Label 0: Type INGRESS-REPLICATION 10.255.162.109
```

```
             Next hop type: Indirect, Next hop index: 0
             Address: 0xa5b8690
             Next-hop reference count: 11
             Protocol next hop: 10.255.162.109
             Indirect next hop: 0x0 - INH Session ID: 0x0
             State: <Active Int Ext>
             Age: 59:50    Metric2: 1
             Validation State: unverified
             Task: mvpn global task
             Announcement bits (3): 0-PIM.vpn1 1-mvpn global task 2-rt-export
```

```
             AS path: I
             Communities: segmented-nh:10.255.162.109:0
```

```
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109/240 (1 entry, 1 announced)
```

```
    *MVPN    Preference: 70
             PMSI: Flags 0x1: Label 0: Type INGRESS-REPLICATION 10.255.162.109
```

```
             Next hop type: Indirect, Next hop index: 0
             Address: 0xa5b8690
             Next-hop reference count: 11
             Protocol next hop: 10.255.162.109
             Indirect next hop: 0x0 - INH Session ID: 0x0
             State: <Active Int Ext>
             Age: 1:00:20    Metric2: 1
             Validation State: unverified
             Task: mvpn global task
             Announcement bits (3): 0-PIM.vpn1 1-mvpn global task 2-rt-export
```

AS path: I  
Communities: segmented-nh:10.255.162.109:0

**Meaning** The output indicates the route table for the segmented type-3 traffic generated from ABR1.

*Verifying the Route Table for Segmented Type-4 Traffic Received from Device ABR1  
Toward PE1 Router*

**Purpose** Verify the route table for segmented type-4 traffic received from Device ABR1.

**Action** From operational mode, run the **show route table vpn1.mvpn.0 match-prefix 4:\* detail** command.

```
user@PE1> show route table vpn1.mvpn.0 match-prefix 4:* detail
```

```
vpn1.mvpn.0: 19 destinations, 22 routes (19 active, 3 holddown, 0 hidden)
4:3:10.255.162.109:100:32:172.16.1.2:32:20192.0.2.2:10.255.162.109:10.255.162.100/240
(1 entry, 1 announced)
    *BGP      Preference: 170/-101
              PMSI: Flags 0x0: Label 300320: Type INGRESS-REPLICATION
10.255.162.100
    Next hop type: Indirect, Next hop index: 0
    Address: 0xa5d11d0
    Next-hop reference count: 24
    Source: 10.255.162.100
    Protocol next hop: 10.255.162.100
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Secondary Active Int Ext>
    Local AS: 65550 Peer AS: 65550
    Age: 1:00:29    Metric2: 2
    Validation State: unverified
    Task: BGP_65550.10.255.162.100
    Announcement bits (2): 0-PIM.vpn1 1-mvpn global task
    AS path: I
    Communities: target:10.255.162.109:0
    Import Accepted
    Localpref: 100
    Router ID: 10.255.162.100
    Primary Routing Table bgp.mvpn.0

4:3:10.255.162.109:100:32:172.16.1.2:32:20192.0.2.1:10.255.162.109:10.255.162.100/240
(1 entry, 1 announced)
    *BGP      Preference: 170/-101
              PMSI: Flags 0x0: Label 300352: Type INGRESS-REPLICATION
10.255.162.100
    Next hop type: Indirect, Next hop index: 0
    Address: 0xa5d11d0
    Next-hop reference count: 24
    Source: 10.255.162.100
    Protocol next hop: 10.255.162.100
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Secondary Active Int Ext>
    Local AS: 65550 Peer AS: 65550
    Age: 59:59     Metric2: 2
    Validation State: unverified
    Task: BGP_65550.10.255.162.100
    Announcement bits (2): 0-PIM.vpn1 1-mvpn global task
    AS path: I
    Communities: target:10.255.162.109:0
    Import Accepted
    Localpref: 100
    Router ID: 10.255.162.100
    Primary Routing Table bgp.mvpn.0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109:10.255.162.100/240
(1 entry, 1 announced)
    *BGP      Preference: 170/-101
              PMSI: Flags 0x0: Label 300336: Type INGRESS-REPLICATION
10.255.162.100
    Next hop type: Indirect, Next hop index: 0
```

```

Address: 0xa5d11d0
Next-hop reference count: 24
Source: 10.255.162.100
Protocol next hop: 10.255.162.100
Indirect next hop: 0x0 - INH Session ID: 0x0
State: <Secondary Active Int Ext>
Local AS: 65550 Peer AS: 65550
Age: 1:00:29 Metric2: 2
Validation State: unverified
Task: BGP_65550.10.255.162.100
Announcement bits (2): 0-PIM.vpn1 1-mvpn global task
AS path: I
Communities: target:10.255.162.109:0
Import Accepted
Localpref: 100
Router ID: 10.255.162.100
Primary Routing Table bgp.mvpn.0

```

**Meaning** The output shows the route table for segmented type-4 traffic received from device ABR1.

#### *Verifying the LDP Traffic Statistics*

**Purpose** Verify the LDP traffic statistics of Device PE1.

**Action** From operational mode, run the **show ldp traffic-statistics** command.

```
user@PE1> show ldp traffic-statistics
```

INET FEC Statistics:

FEC	Type	Packets	Bytes	Shared
10.255.162.100/32	Transit	0	0	No
	Ingress	112882983	33864894900	No
10.255.162.102/32	Transit	0	0	No
	Ingress	3884115	1165234500	No
10.255.162.104/32	Transit	0	0	No
	Ingress	3884115	1165234500	No
10.255.162.107/32	Transit	0	0	No
	Ingress	0	0	No
10.255.162.117/32	Transit	0	0	No
	Ingress	0	0	No
10.255.162.119/32	Transit	0	0	No
	Ingress	0	0	No
198.51.100.19/24	Transit	0	0	No
	Ingress	0	0	No
198.51.100.17/24	Transit	0	0	No
	Ingress	3884115	1165234500	No

**Meaning** The output shows the LDP traffic statistics.

## Configuring ABR1

### 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 ABR1:

1. Configure the interfaces.

```
[edit interfaces]
```

```
user@ABR1# set ge-1/0/4 unit 0 family inet address 203.0.113.12/24
```

```
user@ABR1# set ge-1/0/4 unit 0 family iso
```

```
user@ABR1# set ge-1/0/4 unit 0 family inet6 address ::203.0.113.12/120
```

```
user@ABR1# set ge-1/0/4 unit 0 family mpls
```

```
user@ABR1# set ge-1/1/3 unit 0 family inet address 203.0.113.5/24
```

```
user@ABR1# set ge-1/1/3 unit 0 family iso
```

```
user@ABR1# set ge-1/1/3 unit 0 family inet6 address ::203.0.113.5/120
```

```
user@ABR1# set ge-1/1/3 unit 0 family mpls
```

```
user@ABR1# set ge-1/1/9 unit 0 family inet address 203.0.113.4/24
```

```
user@ABR1# set ge-1/1/9 unit 0 family iso
```

```
user@ABR1# set ge-1/1/9 unit 0 family inet6 address ::203.0.113.4/120
```

```
user@ABR1# set ge-1/1/9 unit 0 family mpls
```

```
user@ABR1# set lo0 unit 0 family inet address 203.0.113.0/24
```

```
user@ABR1# set lo0 unit 0 family inet address 10.255.162.100/32 primary
```

2. Configure the autonomous system number.

```
[edit routing-options]
```

```
user@ABR1# set autonomous-system 65550
```

3. Disable RSVP on the management interface and enable RSVP on the interfaces.

```
[edit protocols rsvp]
```

```
user@ABR1# set interface fxp0.0 disable
```

```
user@ABR1# set interface ge-1/1/9.0
```

```
user@ABR1# set interface ge-1/0/4.0
```

```
user@ABR1# set interface ge-1/1/3.0
```

```
user@ABR1# set interface lo0.0
```

```
user@ABR1# set interface all
```

4. Configure MPLS IPv6 tunneling.

```
[edit protocols mpls]
```

```
user@ABR1# set ipv6-tunneling
```

5. Configure MPLS on the interfaces.

```
[edit protocols mpls]
```

```
user@ABR1# set interface fxp0.0 disable
```

```

user@ABR1# set interface ge-1/1/9.0
user@ABR1# set interface ge-1/0/4.0
user@ABR1# set interface ge-1/1/3.0
user@ABR1# set interface lo0.0
user@ABR1# set interface all

```

6. Configure the BGP protocol.

```

[edit protocols bgp]
user@ABR1# set group IBGP_1 type internal
user@ABR1# set group IBGP_1 local-address 10.255.162.100
user@ABR1# set group IBGP_1 family inet any
user@ABR1# set group IBGP_1 family inet-vpn unicast
user@ABR1# set group IBGP_1 family inet-vpn multicast
user@ABR1# set group IBGP_1 family inet6 any
user@ABR1# set group IBGP_1 family inet6-vpn unicast
user@ABR1# set group IBGP_1 family inet-mvpn signaling
user@ABR1# set group IBGP_1 family inet6-mvpn signaling
user@ABR1# set group IBGP_1 family inet-mdt signaling
user@ABR1# set group IBGP_1 cluster 0.0.0.1
user@ABR1# set group IBGP_1 neighbor 10.255.162.109

user@ABR1# set group IBGP_0 type internal
user@ABR1# set group IBGP_0 local-address 10.255.162.100
user@ABR1# set group IBGP_0 family inet any
user@ABR1# set group IBGP_0 family inet-vpn unicast
user@ABR1# set group IBGP_0 family inet-vpn multicast
user@ABR1# set group IBGP_0 family inet6 any
user@ABR1# set group IBGP_0 family inet6-vpn unicast
user@ABR1# set group IBGP_0 family inet-mvpn signaling
user@ABR1# set group IBGP_0 family inet6-mvpn signaling
user@ABR1# set group IBGP_0 family inet-mdt signaling
user@ABR1# set group IBGP_0 neighbor 10.255.162.117
user@ABR1# set group IBGP_0 neighbor 10.255.162.107

```

7. Configure OSPF traffic engineering attributes and enable OSPF on the interfaces.

```

[edit protocols ospf]
user@ABR1# set traffic-engineering
user@ABR1# set area 0.0.0.1 interface fxp0.0 disable
user@ABR1# set area 0.0.0.1 interface ge-1/1/9.0
user@ABR1# set area 0.0.0.0 interface ge-1/0/4.0
user@ABR1# set area 0.0.0.0 interface ge-1/1/3.0
user@ABR1# set area 0.0.0.0 interface lo0.0

```

8. Enable LDP on all the interfaces and advertise P2MP capability to peers.

```

[edit protocols ldp]
user@ABR1# set interface all
user@ABR1# set p2mp

```

9. Configure PIM on the interfaces.

```

[edit protocols pim]

```

```

user@ABR1# set interface all
user@ABR1# set interface fxp0.0 disable
user@ABR1# set interface lo0.0

```

10. Configure the tunnels of the inter-region template for a specific region or all regions.

```

[edit protocols mvpn inter-region-template]
user@ABR1# set template template_1 region IBGP_0 rsvp-te
label-switched-path-template default-template
user@ABR1# set template template_2 region IBGP_0 ldp-p2mp
user@ABR1# set template template_3 region IBGP_0 ingress-replication
create-new-ucast-tunnel
user@ABR1# set template template_3 region IBGP_0 ingress-replication
label-switched-path label-switched-path-template default-template
user@ABR1# set template template_4 all-regions incoming
user@ABR1# set template template_5 region IBGP_0 rsvp-te static-lsp ABR1_to_ABR3

```

11. Configure the routing instance type, route distinguisher, inter-region template of the provider tunnel, and VRF target community, and advertise a single VPN label for all the routes in the VRF for the routing instance.

```

[edit routing-instances]
user@ABR1# set vpn1 instance-type vrf
user@ABR1# set vpn1 route-distinguisher 10.255.162.100:100
user@ABR1# set vpn1 provider-tunnel inter-region template template_1
user@ABR1# set vpn1 vrf-target target:123:1
user@ABR1# set vpn1 vrf-table-label

```

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

```

user@ABR1# show interfaces
ge-1/0/4 {
  unit 0 {
    family inet {
      address 203.0.113.12/24;
    }
    family iso;
    family inet6 {
      address ::203.0.113.12/120;
    }
    family mpls;
  }
}
ge-1/1/3 {
  unit 0 {
    family inet {
      address 203.0.113.5/24;
    }
    family iso;
    family inet6 {

```



```

        address ::203.0.113.5/120;
    }
    family mpls;
}
}
ge-1/1/9 {
    unit 0 {
        family inet {
            address 203.0.113.4/24;
        }
        family iso;
        family inet6 {
            address ::203.0.113.4/120;
        }
        family mpls;
    }
}
lo0 {
    unit 201 {
        family inet {
            address 203.0.113.0/24;
            address 10.255.162.100/32 {
                primary;
            }
        }
    }
}
}

user@ABR1# show protocols
rsvp {
    interface fxp0.0 {
        disable;
    }
    interface ge-1/1/9.0;
    interface ge-1/0/4.0;
    interface ge-1/1/3.0;
    interface lo0.0;
    interface all;
}
mpls {
    ipv6-tunneling;
    interface fxp0.0 {
        disable;
    }
    interface ge-1/1/9.0;
    interface ge-1/0/4.0;
    interface ge-1/1/3.0;
    interface lo0.0;
    interface all;
}
bgp {
    group IBGP_1 {
        type internal;
        local-address 10.255.162.100;
        family inet {
            any;

```

```
    }
    family inet-vpn {
        unicast;
        multicast;
    }
    family inet6 {
        any;
    }
    family inet6-vpn {
        unicast;
    }
    family inet-mvpn {
        signaling;
    }
    family inet6-mvpn {
        signaling;
    }
    family inet-mdt {
        signaling;
    }
    cluster 0.0.0.1;
    neighbor 10.255.162.109;
}
group IBGP_0 {
    type internal;
    local-address 10.255.162.100;
    family inet {
        any;
    }
    family inet-vpn {
        unicast;
        multicast;
    }
    family inet6 {
        any;
    }
    family inet6-vpn {
        unicast;
    }
    family inet-mvpn {
        signaling;
    }
    family inet6-mvpn {
        signaling;
    }
    family inet-mdt {
        signaling;
    }
    neighbor 10.255.162.117;
    neighbor 10.255.162.107;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.1 {
        interface fxp0.0 {
```

```

        disable;
    }
    interface ge-1/1/9.0;
}
area 0.0.0.0 {
    interface ge-1/0/4.0;
    interface ge-1/1/3.0;
    interface lo0.0;
}
}
ldp {
    interface all;
    p2mp;
}
pim {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
mvpn {
    inter-region-template {
        template template_1 {
            region IBGP_0 {
                rsvp-te {
                    label-switched-path-template {
                        default-template;
                    }
                }
            }
        }
        template template_2 {
            region IBGP_0 {
                ldp-p2mp;
            }
        }
        template template_3 {
            region IBGP_0 {
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template {
                            default-template;
                        }
                    }
                }
            }
        }
        template template_4 {
            all-regions {
                incoming;
            }
        }
        template template_5 {
            region IBGP_0 {

```

```
        rsvp-te {
            static-lsp ABR1_to_ABR3;
        }
    }
}

user@ABR1# show routing-instances
vpn1 {
    instance-type vrf;
    route-distinguisher 10.255.162.100:100;
    provider-tunnel {
        inter-region {
            template template_1;
        }
    }
    vrf-target target:123:1;
    vrf-table-label;
}

user@ABR1# show routing-options
autonomous-system 65550;
```

---

### Verification

Confirm that the configuration is working properly.

***Verifying the Segmented Type-3 Traffic Received from the PE1 Router on ABR1 with the Tunnel-Type as IR***

**Purpose** Display the segmented type-3 traffic received from the PE1 router on ABR1 with the tunnel-type as IR.

**Action** From operational mode, run the **show route table vpn1.mvpn.0 match-prefix 3:\* detail** command.

```
user@ABR1> show route table vpn1.mvpn.0 match-prefix 3:* detail
```

```
vpn1.mvpn.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
```

```
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.2:10.255.162.109/240 (1 entry, 1 announced)
```

```
*BGP      Preference: 170/-101
```

```
PMSI: Flags 0x1: Label 0: Type INGRESS-REPLICATION 10.255.162.109
```

```
Next hop type: Indirect, Next hop index: 0
```

```
Address: 0xa5cddb0
```

```
Next-hop reference count: 24
```

```
Source: 10.255.162.109
```

```
Protocol next hop: 10.255.162.109
```

```
Indirect next hop: 0x0 - INH Session ID: 0x0
```

```
State: <Secondary Active Int Ext>
```

```
Local AS: 65550 Peer AS: 65550
```

```
Age: 1:02:45 Metric2: 2
```

```
Validation State: unverified
```

```
Task: BGP_65550.10.255.162.109
```

```
Announcement bits (1): 0-mvpn global task
```

```
AS path: I
```

```
Communities: target:123:1 segmented-nh:10.255.162.109:0
```

```
Import Accepted
```

```
Localpref: 100
```

```
Router ID: 10.255.162.109
```

```
Primary Routing Table bgp.mvpn.0
```

```
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.1:10.255.162.109/240 (1 entry, 1 announced)
```

```
*BGP      Preference: 170/-101
```

```
PMSI: Flags 0x1: Label 0: Type INGRESS-REPLICATION 10.255.162.109
```

```
Next hop type: Indirect, Next hop index: 0
```

```
Address: 0xa5cddb0
```

```
Next-hop reference count: 24
```

```
Source: 10.255.162.109
```

```
Protocol next hop: 10.255.162.109
```

```
Indirect next hop: 0x0 - INH Session ID: 0x0
```

```
State: <Secondary Active Int Ext>
```

```
Local AS: 65550 Peer AS: 65550
```

```
Age: 1:02:15 Metric2: 2
```

```
Validation State: unverified
```

```
Task: BGP_65550.10.255.162.109
```

```
Announcement bits (1): 0-mvpn global task
```

```
AS path: I
```

```
Communities: target:123:1 segmented-nh:10.255.162.109:0
```

```
Import Accepted
```

```
Localpref: 100
```

```
Router ID: 10.255.162.109
```

```
Primary Routing Table bgp.mvpn.0
```

```
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109/240 (1 entry, 1 announced)
```

```
*BGP      Preference: 170/-101
```

```
PMSI: Flags 0x1: Label 0: Type INGRESS-REPLICATION 10.255.162.109
```

```

Next hop type: Indirect, Next hop index: 0
Address: 0xa5cddb0
Next-hop reference count: 24
Source: 10.255.162.109
Protocol next hop: 10.255.162.109
Indirect next hop: 0x0 - INH Session ID: 0x0
State: <Secondary Active Int Ext>
Local AS: 65550 Peer AS: 65550
Age: 1:02:45 Metric2: 2
Validation State: unverified
Task: BGP_65550.10.255.162.109
Announcement bits (1): 0-mvpn global task
AS path: I
Communities: target:123:1 segmented-nh:10.255.162.109:0
Import Accepted
Localpref: 100
Router ID: 10.255.162.109
Primary Routing Table bgp.mvpn.0

```

**Meaning** The output shows the segmented type-3 traffic received from PE1 with the tunnel-type as IR.

### Configuring ABR2

**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 ABR2:

1. Configure the interfaces.

```
[edit interfaces]
```

```

user@ABR2# set ge-1/0/3 unit 0 family inet address 203.0.113.11/24
user@ABR2# set ge-1/0/3 unit 0 family iso
user@ABR2# set ge-1/0/3 unit 0 family inet6 address ::203.0.113.11/120
user@ABR2# set ge-1/0/3 unit 0 family mpls

```

```

user@ABR2# set ge-1/1/4 unit 0 family inet address 203.0.113.10/24
user@ABR2# set ge-1/1/4 unit 0 family iso
user@ABR2# set ge-1/1/4 unit 0 family inet6 address ::203.0.113.10/120
user@ABR2# set ge-1/1/4 unit 0 family mpls

```

```

user@ABR2# set ge-1/1/10 unit 1 family inet address 192.0.2.2/24
user@ABR2# set ge-1/1/10 unit 1 family inet6 address ::192.0.2.2/120
user@ABR2# set ge-1/1/10 unit 1 family mpls

```

2. Configure the autonomous system number.

```
[edit routing-options]
```

```
user@ABR2# set autonomous-system 65550
```

3. Disable RSVP on the management interface and enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@ABR2# set interface fxp0.0 disable
user@ABR2# set interface lo0.0
user@ABR2# set interface all
```

4. Enable MPLS IPv6 tunneling.

```
[edit protocols mpls]
user@ABR2# set ipv6-tunneling
```

5. Disable MPLS on the management interface and enable RSVP on the interfaces.

```
[edit protocols mpls]
user@ABR2# set interface fxp0.0 disable
user@ABR2# set interface lo0.0
user@ABR2# set interface all
```

6. Configure the BGP protocol.

```
[edit protocols bgp]
user@ABR2# set group IBGP_2 type internal
user@ABR2# set group IBGP_2 local-address 10.255.162.117
user@ABR2# set group IBGP_2 family inet any
user@ABR2# set group IBGP_2 family inet-vpn unicast
user@ABR2# set group IBGP_2 family inet-vpn multicast
user@ABR2# set group IBGP_2 family inet6 any
user@ABR2# set group IBGP_2 family inet6-vpn unicast
user@ABR2# set group IBGP_2 family inet6-mvpn signaling
user@ABR2# set group IBGP_2 family inet6-mvpn signaling
user@ABR2# set group IBGP_2 family inet6-mdt signaling
user@ABR2# set group IBGP_2 cluster 0.0.0.2
user@ABR2# set group IBGP_2 neighbor 10.255.162.104
user@ABR2# set group IBGP_2 neighbor 198.51.100.17
```

```
user@ABR2# set group IBGP_0 type internal
user@ABR2# set group IBGP_0 local-address 10.255.162.117
user@ABR2# set group IBGP_0 family inet any
user@ABR2# set group IBGP_0 family inet-vpn unicast
user@ABR2# set group IBGP_0 family inet-vpn multicast
user@ABR2# set group IBGP_0 family inet6 any
user@ABR2# set group IBGP_0 family inet6-vpn unicast
user@ABR2# set group IBGP_0 family inet6-mvpn signaling
user@ABR2# set group IBGP_0 family inet6-mvpn signaling
user@ABR2# set group IBGP_0 family inet6-mdt signaling
user@ABR2# set group IBGP_0 neighbor 10.255.162.100
user@ABR2# set group IBGP_0 neighbor 10.255.162.107
```

7. Configure OSPF traffic engineering attributes, and disable OSPF on the management interface and enable OSPF on the interfaces.

```
[edit protocols ospf]
user@ABR2# set traffic-engineering
user@ABR2# set area 0.0.0.0 interface fxp0.0 disable
```

```

user@ABR2# set area 0.0.0.0 interface ge-1/0/3.0
user@ABR2# set area 0.0.0.0 interface ge-1/1/4.0
user@ABR2# set area 0.0.0.0 interface lo0.0
user@ABR2# set area 0.0.0.2 interface ge-1/1/10.1

```

8. Enable LDP on all the interfaces and advertise P2MP capability to peers.

```

[edit protocols ldp]
user@ABR2# set interface all
user@ABR2# set p2mp

```

9. Configure PIM on the interfaces.

```

[edit protocols pim]
user@ABR2# set interface fxp0.0 all
user@ABR2# set interface fxp0.0 disable
user@ABR2# set interface lo0.0

```

10. Configure the tunnels of the inter-region template for a specific region or all regions.

```

[edit protocols mvpn inter-region-template]
user@ABR2# set template template_1 region IBGP_2 rsvp-te
label-switched-path-template default-template
user@ABR2# set template template_2 region IBGP_2 ldp-p2mp
user@ABR2# set template template_3 region IBGP_2 ingress-replication
create-new-ucast-tunnel
user@ABR2# set template template_3 region IBGP_2 ingress-replication
label-switched-path label-switched-path-template default-template
user@ABR2# set template template_4 all-regions incoming
user@ABR2# set template template_5 region IBGP_2 rsvp-te static-lsp
ABR2_to_PE2_3

```

11. Configure the routing instance type, route distinguisher, inter-region template of the provider tunnel, and VRF target community, and advertise a single VPN label for all the routes in the VRF for the routing instance.

```

[edit routing-instances]
user@ABR2# set vpn1 instance-type vrf
user@ABR2# set vpn1 route-distinguisher 10.255.162.117:100
user@ABR2# set vpn1 provider-tunnel inter-region template template_1
user@ABR2# set vpn1 vrf-target target:123:1
user@ABR2# set vpn1 vrf-table-label

```

**Results**

```

user@ABR2# show interfaces
ge-1/0/3 {
  unit 0 {
    family inet {
      address 203.0.113.11/24;
    }
    family iso;
    family inet6 {
      address ::203.0.113.11/120;
    }
  }
}

```



```

        family mpls;
    }
}
ge-1/1/4 {
    unit 0 {
        family inet {
            address 203.0.113.10/24;
        }
        family iso;
        family inet6 {
            address ::203.0.113.10/120;
        }
        family mpls;
    }
}
ge-1/1/10 {
    unit 1 {
        family inet {
            address 192.0.2.2/24;
        }
        family inet6 {
            address ::192.0.2.2/120;
        }
        family mpls;
    }
}
lo0 {
    unit 201 {
        family inet {
            address 203.0.113.0/24;
            address 10.255.162.117/32 {
                primary;
            }
        }
    }
}
}

user@ABR2# show protocols
rsvp {
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
    interface all;
}
mpls {
    ipv6-tunneling;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
    interface all;
}
bgp {
    group IBGP_2 {
        type internal;
    }
}

```

```
local-address 10.255.162.117;
family inet {
    any;
}
family inet-vpn {
    unicast;
    multicast;
}
family inet6 {
    any;
}
family inet6-vpn {
    unicast;
}
family inet-mvpn {
    signaling;
}
family inet6-mvpn {
    signaling;
}
family inet-mdt {
    signaling;
}
cluster 0.0.0.2;
neighbor 10.255.162.104;
neighbor 198.51.100.17;
}
group IBGP_0 {
    type internal;
    local-address 10.255.162.117;
    family inet {
        any;
    }
    family inet-vpn {
        unicast;
        multicast;
    }
    family inet6 {
        any;
    }
    family inet6-vpn {
        unicast;
    }
    family inet-mvpn {
        signaling;
    }
    family inet6-mvpn {
        signaling;
    }
    family inet6-mdt {
        signaling;
    }
    neighbor 10.255.162.100;
    neighbor 10.255.162.107;
}
ospf {
```

```

traffic-engineering;
area 0.0.0.0 {
    interface fxp0.0 {
        disable;
    }
    interface ge-1/0/3.0;
    interface ge-1/1/4.0;
    interface lo0.0;
}
area 0.0.0.2 {
    interface ge-1/1/10.1;
}
}
ldp {
    interface all;
    p2mp;
}
pim {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
mvpn {
    inter-region-template {
        template template_1 {
            region IBGP_2 {
                rsvp-te {
                    label-switched-path-template {
                        default-template;
                    }
                }
            }
        }
        template template_2 {
            region IBGP_2 {
                ldp-p2mp;
            }
        }
        template template_3 {
            region IBGP_2 {
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template {
                            default-template;
                        }
                    }
                }
            }
        }
        template template_4 {
            all-regions {
                incoming;
            }
        }
    }
}

```

```
    }
    template template_5 {
      region IBGP_2 {
        rsvp-te {
          static-lsp ABR2_to_PE2_3;
        }
      }
    }
  }
}

user@ABR2# show routing-instances
vpn1 {
  instance-type vrf;
  route-distinguisher 10.255.162.100:100;
  provider-tunnel {
    inter-region {
      template template_1;
    }
  }
  vrf-target target:123:1;
  vrf-table-label;
}

user@ABR2# show routing-options
autonomous-system 65550;
```

---

## Verification

Confirm that the configuration is working properly.

- [Verifying Segmented Type-3 Received from ABR2 on page 520](#)
- [Verifying Type-4 Received from Egress PE2 and PE4 and Locally Triggered Type-4 Toward Ingress ABR2 on page 522](#)
- [Verifying the Statistics of MPLS LSP on page 526](#)

### *Verifying Segmented Type-3 Received from ABR2*

**Purpose** Display the segmented Type-3 received from ABR2 where the tunnel type is RSVP-TE.

**Action** From operational mode, enter the `show route table vpn1.mvpn match-prefix 3:* detail` command.

```
user@ABR2> show route table vpn1.mvpn match-prefix 3:* detail
```

```
vpn1.mvpn.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.2:10.255.162.109/240 (1 entry, 1
announced)
```

```
    *BGP      Preference: 170/-101
              PMSI: Flags 0x1: Label 0: RSVP-TE:
Session_13[10.255.162.100:0:6500:10.255.162.100]
              Next hop type: Indirect, Next hop index: 0
              Address: Oxa5bd650
              Next-hop reference count: 24
              Source: 10.255.162.100
              Protocol next hop: 10.255.162.109
              Indirect next hop: 0x0 - INH Session ID: 0x0
              State: <Secondary Active Int Ext>
              Local AS: 65550 Peer AS: 65550
              Age: 1:10:55    Metric2: 1
              Validation State: unverified
              Task: BGP_65550.10.255.162.100
              Announcement bits (1): 0-mvpn global task
              AS path: I (Originator)
              Cluster list: 0.0.0.1
              Originator ID: 10.255.162.109
              Communities: target:123:1 segmented-nh:10.255.162.100:0
              Import Accepted
              Localpref: 100
              Router ID: 10.255.162.100
              Primary Routing Table bgp.mvpn.0
```

```
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.1:10.255.162.109/240 (1 entry, 1
announced)
```

```
    *BGP      Preference: 170/-101
              PMSI: Flags 0x1: Label 0: RSVP-TE:
Session_13[10.255.162.100:0:6504:10.255.162.100]
              Next hop type: Indirect, Next hop index: 0
              Address: Oxa5bd650
              Next-hop reference count: 24
              Source: 10.255.162.100
              Protocol next hop: 10.255.162.109
              Indirect next hop: 0x0 - INH Session ID: 0x0
              State: <Secondary Active Int Ext>
              Local AS: 65550 Peer AS: 65550
              Age: 1:10:25    Metric2: 1
              Validation State: unverified
              Task: BGP_65550.10.255.162.100
              Announcement bits (1): 0-mvpn global task
              AS path: I (Originator)
              Cluster list: 0.0.0.1
              Originator ID: 10.255.162.109
              Communities: target:123:1 segmented-nh:10.255.162.100:0
              Import Accepted
              Localpref: 100
              Router ID: 10.255.162.100
              Primary Routing Table bgp.mvpn.0
```

```
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109/240 (1 entry, 1
announced)
```

```
*BGP      Preference: 170/-101
          PMSI: Flags 0x1: Label 0: RSVP-TE:
Session_13[10.255.162.100:0:6502:10.255.162.100]
          Next hop type: Indirect, Next hop index: 0
          Address: 0xa5bd650
          Next-hop reference count: 24
          Source: 10.255.162.100
          Protocol next hop: 10.255.162.109
          Indirect next hop: 0x0 - INH Session ID: 0x0
          State: <Secondary Active Int Ext>
          Local AS: 65550 Peer AS: 65550
          Age: 1:10:55    Metric2: 1
          Validation State: unverified
          Task: BGP_65550.10.255.162.100
          Announcement bits (1): 0-mvpn global task
          AS path: I (Originator)
          Cluster list: 0.0.0.1
          Originator ID: 10.255.162.109
          Communities: target:123:1 segmented-nh:10.255.162.100:0
          Import Accepted
          Localpref: 100
          Router ID: 10.255.162.100
          Primary Routing Table bgp.mvpn.0
```

**Meaning** The output displays the segmented Type-3 traffic received from ABR2 where the tunnel type is RSVP-TE.

***Verifying Type-4 Received from Egress PE2 and PE4 and Locally Triggered Type-4 Toward Ingress ABR2***

**Purpose** Display the type-4 received from the egress PE2 and PE4 and locally triggered type-4 toward ingress ABR2.

**Action** From operational mode, enter the **show route table vpn1.mvpn match-prefix 4:\* detail** command.

```
user@ABR2> show route table vpn1.mvpn match-prefix 4:* detail
```

```
vpn1.mvpn.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.2:10.255.162.109:10.255.162.104/240
(1 entry, 1 announced)
```

```
*BGP      Preference: 170/-101
          Next hop type: Indirect, Next hop index: 0
          Address: 0xa5d1720
          Next-hop reference count: 21
          Source: 10.255.162.104
          Protocol next hop: 10.255.162.104
          Indirect next hop: 0x0 - INH Session ID: 0x0
          State: <Secondary Active Int Ext>
          Local AS: 65550 Peer AS: 65550
          Age: 1:11:05   Metric2: 2
          Validation State: unverified
          Task: BGP_65550.10.255.162.104
          Announcement bits (1): 0-mvpn global task
          AS path: I
          Communities: target:10.255.162.117:0
          Import Accepted
          Localpref: 100
          Router ID: 10.255.162.104
          Primary Routing Table bgp.mvpn.0
```

```
4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.2:10.255.162.109:10.255.162.117/240
(1 entry, 1 announced)
```

```
*MVPN      Preference: 70
          Next hop type: Indirect, Next hop index: 0
          Address: 0xa5d31f0
          Next-hop reference count: 11
          Protocol next hop: 10.255.162.117
          Indirect next hop: 0x0 - INH Session ID: 0x0
          State: Active Int Ext
          Age: 1:11:04   Metric2: 1
          Validation State: unverified
          Task: mvpn global task
          Announcement bits (2): 0-mvpn global task 1-rt-export
          AS path: I
          Communities: target:10.255.162.100:0
```

```
4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.2:10.255.162.109:198.51.100.17/240
(1 entry, 1 announced)
```

```
*BGP      Preference: 170/-101
          Next hop type: Indirect, Next hop index: 0
          Address: 0xa5cb0f0
          Next-hop reference count: 21
          Source: 198.51.100.17
          Protocol next hop: 198.51.100.17
          Indirect next hop: 0x0 - INH Session ID: 0x0
          State: <Secondary Active Int Ext>
          Local AS: 65550 Peer AS: 65550
          Age: 1:11:05   Metric2: 3
          Validation State: unverified
          Task: BGP_65550.198.51.100.17
          Announcement bits (1): 0-mvpn global task
          AS path: I
```

```
Communities: target:10.255.162.117:0
Import Accepted
Localpref: 100
Router ID: 198.51.100.17
Primary Routing Table bgp.mvpn.0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.1:10.255.162.109:10.255.162.104/240
(1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Indirect, Next hop index: 0
        Address: 0xa5d1720
        Next-hop reference count: 21
        Source: 10.255.162.104
        Protocol next hop: 10.255.162.104
        Indirect next hop: 0x0 - INH Session ID: 0x0
        State: <Secondary Active Int Ext>
        Local AS: 65550 Peer AS: 65550
        Age: 1:10:35   Metric2: 2
        Validation State: unverified
        Task: BGP_65550.10.255.162.104
        Announcement bits (1): 0-mvpn global task
        AS path: I
        Communities: target:10.255.162.117:0
        Import Accepted
        Localpref: 100
        Router ID: 10.255.162.104
        Primary Routing Table bgp.mvpn.0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.1:10.255.162.109:10.255.162.117/240
(1 entry, 1 announced)
  *MVPN  Preference: 70
        Next hop type: Indirect, Next hop index: 0
        Address: 0xa5d31f0
        Next-hop reference count: 11
        Protocol next hop: 10.255.162.117
        Indirect next hop: 0x0 - INH Session ID: 0x0
        State: <Active Int Ext>
        Age: 1:10:35   Metric2: 1
        Validation State: unverified
        Task: mvpn global task
        Announcement bits (2): 0-mvpn global task 1-rt-export
        AS path: I
        Communities: target:10.255.162.100:0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.1:10.255.162.109:198.51.100.17/240
(1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Indirect, Next hop index: 0
        Address: 0xa5cb0f0
        Next-hop reference count: 21
        Source: 198.51.100.17
        Protocol next hop: 198.51.100.17
        Indirect next hop: 0x0 - INH Session ID: 0x0
        State: Secondary Active Int Ext
        Local AS: 65550 Peer AS: 65550
        Age: 1:10:35   Metric2: 3
        Validation State: unverified
        Task: BGP_65550.198.51.100.17
        Announcement bits (1): 0-mvpn global task
        AS path: I
        Communities: target:10.255.162.117:0
```



```

Import Accepted
Localpref: 100
Router ID: 198.51.100.17
Primary Routing Table bgp.mvpn.0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109:10.255.162.104/240
(1 entry, 1 announced)
*BGP Preference: 170/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xa5d1720
Next-hop reference count: 21
Source: 10.255.162.104
Protocol next hop: 10.255.162.104
Indirect next hop: 0x0 - INH Session ID: 0x0
State: <Secondary Active Int Ext>
Local AS: 65550 Peer AS: 65550
Age: 1:11:04 Metric2: 2
Validation State: unverified
Task: BGP_65550.10.255.162.104
Announcement bits (1): 0-mvpn global task
AS path: I
Communities: target:10.255.162.117:0
Import Accepted
Localpref: 100
Router ID: 10.255.162.104
Primary Routing Table bgp.mvpn.0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109:10.255.162.117/240
(1 entry, 1 announced)
*MVPN Preference: 70
Next hop type: Indirect, Next hop index: 0
Address: 0xa5d31f0
Next-hop reference count: 11
Protocol next hop: 10.255.162.117
Indirect next hop: 0x0 - INH Session ID: 0x0
State: <Active Int Ext>
Age: 1:11:04 Metric2: 1
Validation State: unverified
Task: mvpn global task
Announcement bits (2): 0-mvpn global task 1-rt-export
AS path: I
Communities: target:10.255.162.100:0

4:3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109:198.51.100.17/240
(1 entry, 1 announced)
*BGP Preference: 170/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xa5cb0f0
Next-hop reference count: 21
Source: 198.51.100.17
Protocol next hop: 198.51.100.17
Indirect next hop: 0x0 - INH Session ID: 0x0
State: <Secondary Active Int Ext>
Local AS: 65550 Peer AS: 65550
Age: 1:11:04 Metric2: 3
Validation State: unverified
Task: BGP_65550.198.51.100.17
Announcement bits (1): 0-mvpn global task
AS path: I
Communities: target:10.255.162.117:0
Import Accepted

```

```
Localpref: 100
Router ID: 198.51.100.17
Primary Routing Table bgp.mvpn.0
```

**Meaning** The output shows that the configured tunnel type on the ABR2 is RSVP-TE. The RSVP tunnel from the ABR1 ends in ABR2 as the egress LSP, and the new LSP is triggered to egress PE2 and PE4.

#### *Verifying the Statistics of MPLS LSP*

**Purpose** Display the statistics of MPLS LSP.

**Action** From operational mode, run the **show mpls lsp statistics** command for Device ABR2.

```
user@ABR2> show mpls lsp statistics
```

Ingress LSP: 6 sessions

To	From	State	Packets	Bytes	LSPname
10.255.162.104	10.255.162.117	Up	0	0	
10.255.162.104:10.255.162.117:100:mv20:vpn1					
10.255.162.104	10.255.162.117	Up	0	0	
10.255.162.104:10.255.162.117:100:mv21:vpn1					
10.255.162.104	10.255.162.117	Up	0	0	
10.255.162.104:10.255.162.117:100:mv22:vpn1					
198.51.100.17	10.255.162.117	Up	0	0	
198.51.100.17:10.255.162.117:100:mv20:vpn1					
198.51.100.17	10.255.162.117	Up	0	0	
198.51.100.17:10.255.162.117:100:mv21:vpn1					
198.51.100.17	10.255.162.117	Up	0	0	
198.51.100.17:10.255.162.117:100:mv22:vpn1					

Total 6 displayed, Up 6, Down 0

Egress LSP: 6 sessions

To	From	State	Packets	Bytes	LSPname
10.255.162.117	10.255.162.100	Up	NA	NA	
10.255.162.117:10.255.162.100:100:mv45:vpn1					
10.255.162.117	10.255.162.100	Up	NA	NA	
10.255.162.117:10.255.162.100:100:mv47:vpn1					
10.255.162.117	10.255.162.100	Up	NA	NA	
10.255.162.117:10.255.162.100:100:mv49:vpn1					
10.255.162.117	10.255.162.104	Up	NA	NA	PE2_1_to_ABR2
10.255.162.117	10.255.162.107	Up	NA	NA	ABR3_to_ABR2
10.255.162.117	198.51.100.17	Up	NA	NA	NA
PE2_3_to_ABR2					

Total 6 displayed, Up 6, Down 0

### Configuring ABR3

**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 ABR3:

1. Configure the interfaces.

```
[edit interfaces]
```

```
user@ABR3# set ge-1/0/8 unit 0 family inet address 203.0.113.6/24
```

```
user@ABR3# set ge-1/0/8 unit 0 family iso
```

```
user@ABR3# set ge-1/0/8 unit 0 family inet6 address ::203.0.113.6/120
```

```
user@ABR3# set ge-1/0/8 unit 0 family mpls
```

```
user@ABR3# set ge-1/1/4 unit 0 family inet address 203.0.113.9/24
```

```
user@ABR3# set ge-1/1/4 unit 0 family iso
```

```
user@ABR3# set ge-1/1/4 unit 0 family inet6 address ::203.0.113.9/120
```

```
user@ABR3# set ge-1/1/4 unit 0 family mpls
```

```
user@ABR3# set ge-1/3/1 unit 0 family inet address 203.0.113.8/24
```

```

user@ABR3# set ge-1/3/1 unit 0 family iso
user@ABR3# set ge-1/3/1 unit 0 family inet6 address ::203.0.113.8/120
user@ABR3# set ge-1/3/1 unit 0 family mpls

user@ABR3# set lo0 unit 0 family inet address 203.0.113.0/24
user@ABR3# set lo0 unit 0 family inet address 10.255.162.107/32 primary

```

2. Configure the autonomous system number.

```

[edit routing-options]
user@ABR3# set autonomous-system 65550

```

3. Configure RSVP on all the interfaces, excluding the management interface.

```

[edit protocols rsvp]
user@ABR3# set interface all
user@ABR3# set interface fxp0.0 disable
user@ABR3# set interface lo0.0

```

4. Configure MPLS IPv6 tunneling, configure the label-switched path, and enable MPLS on all the interfaces, excluding the management interface.

```

[edit protocols mpls]
user@ABR3# set ipv6-tunneling
user@ABR3# set label-switched-path ABR3_to_PE3 from 10.255.162.107
user@ABR3# set label-switched-path ABR3_to_PE3 to 10.255.162.102
user@ABR3# set label-switched-path ABR3_to_PE3 p2mp vpn1
user@ABR3# set label-switched-path ABR3_to_ABR1 from 10.255.162.107
user@ABR3# set label-switched-path ABR3_to_ABR1 to 10.255.162.100
user@ABR3# set label-switched-path ABR3_to_ABR1 p2mp vpn1
user@ABR3# set label-switched-path ABR3_to_ABR2 from 10.255.162.107
user@ABR3# set label-switched-path ABR3_to_ABR2 to 10.255.162.117
user@ABR3# set label-switched-path ABR3_to_ABR2 p2mp vpn1
user@ABR3# set interface all
user@ABR3# set interface fxp0.0 disable
user@ABR3# set interface lo0.0

```

5. Configure the BGP protocol.

```

[edit protocols bgp]
user@ABR3# set group IBGP_3 type internal
user@ABR3# set group IBGP_3 local-address 10.255.162.107
user@ABR3# set group IBGP_3 family inet any
user@ABR3# set group IBGP_3 family inet-vpn unicast
user@ABR3# set group IBGP_3 family inet-vpn multicast
user@ABR3# set group IBGP_3 family inet6 any
user@ABR3# set group IBGP_3 family inet6-vpn unicast
user@ABR3# set group IBGP_3 family inet6-mvpn signaling
user@ABR3# set group IBGP_3 family inet6-mvpn signaling
user@ABR3# set group IBGP_3 family inet-mdt signaling
user@ABR3# set group IBGP_3 cluster 0.0.0.3
user@ABR3# set group IBGP_3 neighbor 10.255.162.102

```

```

user@ABR3# set group IBGP_0 type internal

```

```

user@ABR3# set group IBGP_0 local-address 10.255.162.107
user@ABR3# set group IBGP_0 family inet any
user@ABR3# set group IBGP_0 family inet-vpn unicast
user@ABR3# set group IBGP_0 family inet-vpn multicast
user@ABR3# set group IBGP_0 family inet6 any
user@ABR3# set group IBGP_0 family inet6-vpn unicast
user@ABR3# set group IBGP_0 family inet-mvpn signaling
user@ABR3# set group IBGP_0 family inet6-mvpn signaling
user@ABR3# set group IBGP_0 family inet-mdt signaling
user@ABR3# set group IBGP_0 neighbor 10.255.162.100
user@ABR3# set group IBGP_0 neighbor 10.255.162.117

```

6. Configure OSPF traffic engineering attributes, disable OSPF on the management interface, and enable OSPF on the interfaces.

```

[edit protocols ospf]
user@ABR3# set traffic-engineering
user@ABR3# set area 0.0.0.0 interface fxp0.0 disable
user@ABR3# set area 0.0.0.0 interface ge-1/0/8.0
user@ABR3# set area 0.0.0.0 interface ge-1/1/4.0
user@ABR3# set area 0.0.0.0 interface lo0.0
user@ABR3# set area 0.0.0.3 interface ge-1/3/1.0

```

7. Enable LDP on all the interfaces and advertise P2MP capability to peers.

```

[edit protocols ldp]
user@ABR3# set interface all
user@ABR3# set p2mp

```

8. Configure PIM on the interfaces.

```

[edit protocols pim]
user@ABR3# set interface all
user@ABR3# set interface fxp0.0 disable
user@ABR3# set interface lo0.0

```

9. Configure the tunnels of the inter-region template for a specific region or all regions.

```

[edit protocols mvpn inter-region-template]
user@ABR3# set template template_1 region IBGP_3 rsvp-te
  label-switched-path-template default-template
user@ABR3# set template template_2 region IBGP_3 ldp-p2mp
user@ABR3# set template template_3 region IBGP_3 ingress-replication
  create-new-ucast-tunnel
user@ABR3# set template template_3 region IBGP_3 ingress-replication
  label-switched-path label-switched-path-template default-template
user@ABR3# set template template_4 all-regions incoming
user@ABR3# set template template_5 region IBGP_3 rsvp-te static-lsp ABR3_to_PE3

```

10. Configure the routing instance type, route distinguisher, inter-region template of the provider tunnel, and VRF target community, and advertise a single VPN label for all the routes in the VRF for the routing instance.

```

[edit routing-instances]

```

```
user@ABR3# set vpn1 instance-type vrf
user@ABR3# set vpn1 route-distinguisher 10.255.162.107:100
user@ABR3# set vpn1 provider-tunnel inter-region template template_1
user@ABR3# set vpn1 vrf-target target:123:1
user@ABR3# set vpn1 vrf-table-label
```

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

```
user@ABR3# show interfaces
ge-1/0/8 {
  unit 0 {
    family inet {
      address 203.0.113.6/24;
    }
    family iso;
    family inet6 {
      address ::203.0.113.6/120;
    }
    family mpls;
  }
}
ge-1/1/4 {
  unit 0 {
    family inet {
      address 203.0.113.9/24;
    }
    family iso;
    family inet6 {
      address ::203.0.113.9/120;
    }
    family mpls;
  }
}
ge-1/3/1 {
  unit 0 {
    family inet {
      address 203.0.113.8/24;
    }
    family iso;
    family inet6 {
      address ::203.0.113.8/120;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 203.0.113.0/24;
      address 10.255.162.107/32 {
        primary;
      }
    }
  }
}
```

```

    }
  }
}
}
user@ABR3# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0;
}
mpls {
  ipv6-tunneling;
  label-switched-path ABR3_to_PE3 {
    from 10.255.162.107;
    to 10.255.162.102;
    p2mp vpn1;
  }
  label-switched-path ABR3_to_ABR1 {
    from 10.255.162.107;
    to 10.255.162.100;
    p2mp vpn1;
  }
  label-switched-path ABR3_to_ABR2 {
    from 10.255.162.107;
    to 10.255.162.117;
    p2mp vpn1;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0;
}
bgp {
  group IBGP_3 {
    type internal;
    local-address 10.255.162.107;
    family inet {
      any;
    }
    family inet-vpn {
      unicast;
      multicast;
    }
    family inet6 {
      any;
    }
    family inet6-vpn {
      unicast;
    }
    family inet-mvpn {
      signaling;
    }
  }
}

```

```
    family inet6-mvpn {
        signaling;
    }
    family inet-mdt {
        signaling;
    }
    cluster 0.0.0.3;
    neighbor 10.255.162.102;
}
group IBGP_0 {
    type internal;
    local-address 10.255.162.107;
    family inet {
        any;
    }
    family inet-vpn {
        unicast;
        multicast;
    }
    family inet6 {
        any;
    }
    family inet6-vpn {
        unicast;
    }
    family inet-mvpn {
        signaling;
    }
    family inet6-mvpn {
        signaling;
    }
    family inet-mdt {
        signaling;
    }
    neighbor 10.255.162.100;
    neighbor 10.255.162.117;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface ge-1/0/8.0;
        interface ge-1/1/4.0;
        interface lo0.0;
    }
    area 0.0.0.3 {
        interface ge-1/3/1.0;
    }
}
}
ldp {
    interface all;
    p2mp;
}
```



```

pim {
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0;
}
mvpn {
  inter-region-template {
    template template_1 {
      region IBGP_3 {
        rsvp-te {
          label-switched-path-template {
            default-template;
          }
        }
      }
    }
  }
  template template_2 {
    region IBGP_3 {
      ldp-p2mp;
    }
  }
  template template_3 {
    region IBGP_3 {
      ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path {
          label-switched-path-template {
            default-template;
          }
        }
      }
    }
  }
  template template_4 {
    all-regions {
      incoming;
    }
  }
  template template_5 {
    region IBGP_3 {
      rsvp-te {
        static-lsp ABR3_to_PE3_1;
      }
    }
  }
}

user@ABR3# show routing-instances
vpn1 {
  instance-type vrf;
  route-distinguisher 10.255.162.107:100;
  provider-tunnel {
    inter-region {

```

```

        template template_1;
    }
}
vrf-target target:123:1;
vrf-table-label;
}

user@ABR3# show routing-option
autonomous-system 65550;

```

## Verification

Confirm that the configuration is working properly.

### Verifying Segmented Type-3 Received from ABR1 on ABR3

**Purpose** Display the segmented Type-3 received from ABR1 on ABR3 where the tunnel type is RSVP-TE.

**Action** From operational mode, run the `show route table vpn1.mvpn match-prefix 3:* detail` command for Device ABR3.

```

user@ABR3> show route table vpn1.mvpn match-prefix 3:* detail

vpn1.mvpn.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.2:10.255.162.109/240 (1 entry, 1
announced)
    *BGP      Preference: 170/-101
              PMSI: Flags 0x1: Label 0: RSVP-TE:
Session_13[10.255.162.100:0:6500:10.255.162.100]
    Next hop type: Indirect, Next hop index: 0
    Address: 0xa5bd650
    Next-hop reference count: 24
    Source: 10.255.162.100
    Protocol next hop: 10.255.162.109
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Secondary Active Int Ext>
    Local AS: 65550 Peer AS: 65550
    Age: 1:10:55    Metric2: 1
    Validation State: unverified
    Task: BGP_65550.10.255.162.100
    Announcement bits (1): 0-mvpn global task
    AS path: I (Originator)
    Cluster list: 0.0.0.1
    Originator ID: 10.255.162.109
    Communities: target:123:1 segmented-nh:10.255.162.100:0
    Import Accepted
    Localpref: 100
    Router ID: 10.255.162.100
    Primary Routing Table bgp.mvpn.0

3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.1:10.255.162.109/240 (1 entry, 1
announced)
    *BGP      Preference: 170/-101
              PMSI: Flags 0x1: Label 0: RSVP-TE:
Session_13[10.255.162.100:0:6504:10.255.162.100]
    Next hop type: Indirect, Next hop index: 0

```

```

Address: 0xa5bd650
Next-hop reference count: 24
Source: 10.255.162.100
Protocol next hop: 10.255.162.109
Indirect next hop: 0x0 - INH Session ID: 0x0
State: <Secondary Active Int Ext>
Local AS: 65550 Peer AS: 65550
Age: 1:10:25 Metric2: 1
Validation State: unverified
Task: BGP_65550.10.255.162.100
Announcement bits (1): 0-mvpn global task
AS path: I (Originator)
Cluster list: 0.0.0.1
Originator ID: 10.255.162.109
Communities: target:123:1 segmented-nh:10.255.162.100:0
Import Accepted
Localpref: 100
Router ID: 10.255.162.100
Primary Routing Table bgp.mvpn.0

3:10.255.162.109:100:32:172.16.1.2:32:192.0.2.3:10.255.162.109/240 (1 entry, 1
announced)
    *BGP Preference: 170/-101
    PMSI: Flags 0x1: Label 0: RSVP-TE:
Session_13[10.255.162.100:0:6502:10.255.162.100]
    Next hop type: Indirect, Next hop index: 0
    Address: 0xa5bd650
    Next-hop reference count: 24
    Source: 10.255.162.100
    Protocol next hop: 10.255.162.109
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Secondary Active Int Ext>
    Local AS: 65550 Peer AS: 65550
    Age: 1:10:55 Metric2: 1
    Validation State: unverified
    Task: BGP_65550.10.255.162.100
    Announcement bits (1): 0-mvpn global task
    AS path: I (Originator)
    Cluster list: 0.0.0.1
    Originator ID: 10.255.162.109
    Communities: target:123:1 segmented-nh:10.255.162.100:0
    Import Accepted
    Localpref: 100
    Router ID: 10.255.162.100
    Primary Routing Table bgp.mvpn.0

```

**Meaning** The output displays the segmented Type-3 traffic received from ABR1 where the tunnel type is RSVP-TE.

**Related Documentation**

- [Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471](#)
- [Configuring Segmented Inter-Area P2MP LSP on page 476](#)
- [all-regions on page 1042](#)
- [inter-region on page 1082](#)
- [inter-region-template on page 1085](#)

- [inter-region-segmented on page 1083](#)
- [region on page 1122](#)
- [template on page 1141](#)

# Associating Routing Instances with Sender and Receiver Sites

- [Configuring Routing Instances for an MBGP MVPN on page 537](#)
- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 538](#)
- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 540](#)
- [Configuring Internet Multicast Using Ingress Replication Provider Tunnels on page 542](#)
- [Example: Configuring PIM State Limits on page 545](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559](#)
- [Configuring a Selective Provider Tunnel Using Wildcards on page 564](#)
- [Example: Configuring Selective Provider Tunnels Using Wildcards on page 565](#)

## Configuring Routing Instances for an MBGP MVPN

---

To configure MBGP MVPNs, include the **mvpn** statement:

```
mvpn {  
  mvpn-mode (rpt-spt | spt-only);  
  receiver-site;  
  route-target {  
    export-target {  
      target target-community;  
      unicast;  
    }  
    import-target {  
      target {  
        target-value;  
        receiver target-value;  
        sender target-value;  
      }  
      unicast {  
        receiver;  
        sender;  
      }  
    }  
  }  
}
```

```

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

```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

By default an MBGP MVPN routing instance is associated with both the multicast sender and the receiver sites. If you configure the **receiver-site** option, the routing instance is associated with only multicast receiver sites. Configuring the **sender-site** option associates the routing instance with only multicast sender sites.



**NOTE:** When you configure the routing instance for the MBGP MVPN, you must configure MPLS LSPs (either RSVP-signaled or LDP-signaled) between the PE routers of the routing instance to ensure VPN unicast connectivity. Point-to-multipoint LSPs are used for multicast data forwarding only.

## Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs

For MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), the default mode of operation supports only intersite shortest-path trees (SPTs) for customer PIM (C-PIM) join messages. It does not support rendezvous-point trees (RPTs) for C-PIM join messages. The default mode of operation provides advantages, but it requires either that the customer rendezvous point (C-RP) be located on a PE router or that the Multicast Source Discovery Protocol (MSDP) be used between the C-RP and a PE router so that the PE router can learn about active sources advertised by other PE routers.

If the default mode is not suitable for your environment, you can configure RPT-SPT mode (also known as *shared-tree data distribution*), as documented in section 13 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). RPT-SPT mode supports the native PIM model of transmitting (\*G) messages from the receiver to the RP for intersite shared-tree join messages. This means that the type 6 (\*G) routes get transmitted from one PE router to another. In RPT-SPT mode, the shared-tree multicast routes are advertised from an egress PE router to the upstream router connected to the VPN site with the C-RP. The single-forwarder election is performed for the C-RP rather than for the source. The egress PE router takes the upstream hop to advertise the (\*G) and sends the type 6 route toward the upstream PE router. To send the data on the RPT, either inclusive or selective provider tunnels can be used. After the data starts flowing on the RPT, the last-hop router switches to SPT mode, unless you include the **spt-threshold infinity** statements in the configuration.



**NOTE:** The MVPN single-forwarder election follows the rule documented in section 9.1.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). The single-forwarder election winner is based on the following rules:

- If the active unicast route to the source is through the interface, then this route is used to determine the upstream multicast hop (UMH).
- If the active unicast route to the source is a VPN route, MVPN selects the UMH based on the highest IP address in the route import community for the VPN routes, and the local master loopback address for local VRF routes.

The switch to SPT mode is performed by PIM and not by MVPN type 5 and type 6 routes. After the last-hop router switches to SPT mode, the SPT (S,G) join messages follow the same rules as the SPT-only default mode.

The advantage of RPT-SPT mode is that it provides a method for PE routers to discover sources in the multicast VPN when the C-RP is located on the customer site instead of on a PE router. Because the shared C-tree is established between VPN sites, there is no need to run MSDP between the C-RP and the PE routers. RPT-SPT mode also enables egress PE routers to switch to receiving data from the PE connected to the source after the source information is learned, instead of receiving data from the RP.

In Junos OS Release 15.1 and later, in RPT-SPT mode, PIM SSG Joins are created on the egress PE even if no directly-connected receivers are present.



**CAUTION:** When you configure RPT-SPT mode, receivers or sources directly attached to the PE router are not supported. As a workaround, place a CE router between any receiver or source and the PE router.

To configure RPT-SPT mode:

1. Enable shared-tree data distribution:

```
[edit routing-instances routing-instance-name protocols mvpn mvpn-mode]
user@router# set rpt-spt
```

2. Include the **rpt-spt** statement for all VRFs that make up the VPN.

#### Related Documentation

- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 540](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559](#)

## Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

---

For MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), the default mode of operation is shortest path tree only (SPT-only) mode. In SPT-only mode, the active multicast sources are learned through multicast VPN source-active routes. This mode of operation is described in section 14 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt).

In contrast to SPT-only mode, rendezvous point tree (RPT)-SPT mode (also known as shared-tree data distribution) supports the native PIM model of transmitting (\*G) messages from the receiver to the RP for intersite shared-tree join messages.

In SPT-only mode, when a PE router receives a (\*, C-G) join message, the router looks for an active source transmitting data to the customer group. If the PE router has a source-active route for the customer group, the router creates a source tree customer multicast route and sends the route to the PE router connected to the VPN site with the source. The source is determined by MVPN's single-forwarder election. When a receiver sends a (\*G) join message in a VPN site, the (\*G) join message only travels as far as the PE router. After the join message is converted to a type 7 multicast route, which is equivalent to a (S,G) join message, the route is installed with the no-advertise community setting.



**NOTE:** The MVPN single-forwarder election follows the rule documented in section 9.1.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). The single-forwarder election winner is based on the following rules:

- If the active unicast route to the source is through the interface, then this route is used to determine the upstream multicast hop (UMH).
- If the active unicast route to the source is a VPN route, MVPN selects the UMH based on the highest IP address in the route import community for the VPN routes, and the local master loopback address for local VRF routes.

Single-forwarder election guarantees selection of a unique forwarder for a given customer source (C-S). The upstream PE router might differ for the source tree and the shared tree because the election is based on the customer source and C-RP, respectively. Although the single-forwarder election is sufficient for SPT-only mode, the alternative RPT-SPT mode involves procedures to prevent duplicate traffic from being sent on the shared tree and the source tree. These procedures might require administrator-configured parameters to reduce duplicate traffic and reduce blackholes during RPT to SPT switch and the reverse.

In SPT-only mode, when a source is active, PIM creates a register state for the source both on the DR and on the C-RP (or on a PE router that is running Multicast Source Discovery Protocol [MSDP] between itself and the C-RP). After the register states are created, MVPN creates a source-active route. These type 5 source-active routes are installed on all PE routers. When the egress PE router with the (\*G) join message receives



the source-active route, it has two routes that it can combine to produce the (S,G) multicast route. The type 7 route informs the PE router that a receiver is interested in group G. The source active route informs the PE router that a source S is transmitting data to group G. MVPN combines this information to produce a multicast join message and advertises this to the ingress PE router, as determined by the single-forwarder election.

For some service providers, the SPT-only implementation is not ideal because it creates a restriction on C-RP configuration. For a PE router to create customer multicast routes from (\*,C-G) join messages, the router must learn about active sources through MVPN type 5 source-active routes. These source-active routes can be originated only by a PE router. This means that a PE router in the MVPN must learn about all PIM register messages sent to the RP, which is possible only in the following cases:

- The C-RP is colocated on one of the PEs in the MVPN.
- MSDP is run between the C-RP and the VRF instance on one of the PE routers in the MVPN.

If this restriction is not acceptable, providers can use RPT-SPT mode instead of the default SPT-only mode. However, because SPT-only mode does not transmit (\*,G) routes between VPN sites, SPT-only mode has the following advantages over RPT-SPT mode:

- Simplified operations by exchanging and processing only source-tree customer multicast routes among PE routers
- Simplified operations by eliminating the need for the service provider to suppress MVPN transient duplicates during the switch from RPT to SPT
- Less control plane overhead in the service provider space by limiting the type of customer multicast routes exchanged, which results in more scalable deployments
- More stable traffic patterns in the backbone without the traffic shifts involved in the RPT-SPT mode
- Easier maintenance in the service provider space due to less state information

To configure SPT-only mode:

1. Explicitly configure SPT-only mode:

```
[edit routing-instances routing-instance-name protocols mvpn mvpn-mode]
user@router# set spt-only
```

2. Include the **spt-only** statement for all VRFs that make up the VPN.

#### Related Documentation

- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 538](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559](#)

## Configuring Internet Multicast Using Ingress Replication Provider Tunnels

---


The routing instance type **mpls-internet-multicast** uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, enabling a faster path for multicast traffic between sender and receiver routers in large-scale implementations.

The **mpls-internet-multicast** routing instance is a non-forwarding instance used only for control plane procedures; it does not support any interface configurations. Only one **mpls-internet-multicast** routing instance can be defined for a logical system. All multicast and unicast routes used for Internet multicast are associated only with the master instance (inet.0), not with the routing instance.

Each router participating in Internet multicast must be configured with BGP MPLS-based Internet multicast for control plane procedures and with ingress replication for the data provider tunnel, which forms a full mesh of MPLS point-to-point LSPs. The ingress replication tunnel can be selective or inclusive, matching the configuration of the provider tunnel in the routing instance.

The topology consists of routers on the edge of the IP multicast domain that have a set of IP interfaces and a set of MPLS core-facing interfaces, see [Figure 54 on page 543](#). Internet multicast traffic is carried between the IP routers, through the MPLS cloud, using ingress replication tunnels for the data plane and a full-mesh IGBP session for the control plane.

Figure 54: Internet Multicast Topology



The `mpls-internet-multicast` routing instance type is configured for the default master instance on each router to support Internet multicast over MPLS. When using PIM as the multicast protocol, the `mpls-internet-multicast` configuration statement is also included at the `[edit protocols pim]` hierarchy level in the master instance. This creates a pseudo-interface that associates PIM with the `mpls-internet-multicast` routing instance.

When a new destination needs to be added to the ingress replication provider tunnel, the resulting behavior differs depending on how the ingress replication provider tunnel is configured:

- **`create-new-ucast-tunnel`**—When this statement is configured, a new unicast tunnel to the destination is created, and is deleted when the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.
- **`label-switched-path-template (Multicast)`**—When this statement is configured, an LSP template is used for the point-to-multipoint LSP for ingress replication.

**Example: Configure Internet Multicast Using Ingress Replication Tunnels**

This example configures VPN-B with the instance type **mpls-internet-multicast**. This example also uses PIM for the multicast protocol.

1. Configure the routing instance type for VPN-B as **mpls-internet-multicast**:

```
user@host# set routing-instances VPN-B instance-type mpls-internet-multicast
```

2. Configure the ingress replication provider tunnel to create a new unicast tunnel each time an application requests to add a destination:

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
create-new-ucast-tunnel
```

3. Configure the point-to-point LSP to use the default template settings.

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
label-switched-path label-switched-path-template default-template
```

4. Configure the ingress replication provider tunnel to be selective:

```
user@host# set routing-instances VPN-B provider-tunnel selective group
203.0.113.1/24 source 192.168.195.145/32 ingress-replication label-switched-path
```

5. Configure MVPN protocol in the routing instance:

```
user@host# set routing-instances VPN-B protocols mvpn
```

6. Commit the configuration:

```
user@host# commit
```

7. Use show command to verify the instance has been created:

```
user@host# run show mvpn instance VPN-B
MVPN instance:
Legend for provider tunnel I-P-tnl -- inclusive provider tunnel S-P-tnl --
selective provider tunnel
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance : VPN-B
MVPN Mode : SPT-ONLY
Provider tunnel: I-P-tnl:INGRESS-REPLICATION:MPLS Label 18:10.255.245.6
Neighbor          I-P-tnl
10.255.245.2       INGRESS-REPLICATION:MPLS Label 22:10.255.245.2
10.255.245.7       INGRESS-REPLICATION:MPLS Label 19:10.255.245.7
C-mcast IPv4 (S:G) Ptnl          St
192.168.195.145/32:203.0.113.1/24 INGRESS-REPLICATION:MPLS Label
18:10.255.245.6      RM
```

8. Add the **mpls-internet-multicast** configuration statement under the **[edit protocols pim]** hierarchy level in the master instance:

```
user@host# set protocols pim mpls-internet-multicast
```

9. Commit the configuration:

```
user@host# commit
```

10. Use **show ingress-replication mvpn** command to verify configuration settings:

```
user@host# run show ingress-replication mvpn
Ingress Tunnel: mvpn:1
Application: MVPN
Unicast tunnels
  Leaf Address      Tunnel-type      Mode      State
  10.255.245.2      P2P LSP         New       Up
  10.255.245.4      P2P LSP         New       Up
Ingress Tunnel: mvpn:2
Application: MVPN
Unicast tunnels
  Leaf Address      Tunnel-type      Mode      State
  10.255.245.2      P2P LSP         Existing  Up
```

11. Use this if you want to configure the ingress replication provider tunnel to be inclusive:

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
create-new-ucast-tunnel
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
label-switched-path label-switched-path-template default-template
```

12. Use **show mvpn instance** command to verify tunnel is inclusive:

```
user@host# run show mvpn instance VPN-B
MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance : VPN-A
MVPN Mode : SPT-ONLY
Provider tunnel: I-P-tnl:INGRESS-REPLICATION:MPLS Label 18:10.255.245.6
Neighbor      I-P-tnl
10.255.245.2   INGRESS-REPLICATION:MPLS Label 22:10.255.245.2
10.255.245.7   INGRESS-REPLICATION:MPLS Label 19:10.255.245.7
C-mcast IPv4 (S:G) Ptnl      St
192.168.195.145/32:203.0.113.1/24 INGRESS-REPLICATION:MPLS Label
18:10.255.245.6      RM
```

- Related Documentation**
- [create-new-ucast-tunnel on page 1048](#)
  - [ingress-replication on page 1078](#)
  - [mpls-internet-multicast on page 1100](#)

## Example: Configuring PIM State Limits

- [Controlling PIM Resources for Multicast VPNs Overview on page 546](#)
- [Example: Configuring PIM State Limits on page 548](#)

## Controlling PIM Resources for Multicast VPNs Overview

A service provider network must protect itself from potential attacks from misconfigured or misbehaving customer edge (CE) devices and their associated VPN routing and forwarding (VRF) routing instances. Misbehaving CE devices can potentially advertise a large number of multicast routes toward a provider edge (PE) device, thereby consuming memory on the PE device and using other system resources in the network that are reserved for routes belonging to other VPNs.

To protect against potential misbehaving CE devices and VRF routing instances for specific multicast VPNs (MVPNs), you can control the following Protocol Independent Multicast (PIM) resources:

- Limit the number of accepted PIM join messages for any-source groups (\*G) and source-specific groups (S,G).

Note how the device counts the PIM join messages:

- Each (\*G) counts as one group toward the limit.
  - Each (S,G) counts as one group toward the limit.
- Limit the number of PIM register messages received for a specific VRF routing instance. Use this configuration if the device is configured as a rendezvous point (RP) or has the potential to become an RP. When a source in a multicast network becomes active, the source's designated router (DR) encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

Note how the device counts PIM register messages:

- Each unique (S,G) join received by the RP counts as one group toward the configured register messages limit.
  - Periodic register messages sent by the DR for existing or already known (S,G) entries do not count toward the configured register messages limit.
  - Register messages are accepted until either the PIM register limit or the PIM join limit (if configured) is exceeded. Once either limit is reached, any new requests are dropped.
- Limit the number of group-to-RP mappings allowed in a specific VRF routing instance. Use this configuration if the device is configured as an RP or has the potential to become an RP. This configuration can apply to devices configured for automatic RP announce and discovery (Auto-RP) or as a PIM bootstrap router. Every multicast device within a PIM domain must be able to map a particular multicast group address to the same RP. Both Auto-RP and the bootstrap router functionality are the mechanisms used to learn the set of group-to-RP mappings. Auto-RP is typically used in a PIM dense-mode deployment, and the bootstrap router is typically used in a PIM sparse-mode deployment.



**NOTE:** The group-to-RP mappings limit does not apply to static RP or embedded RP configurations.

---

Some important things to note about how the device counts group-to-RP mappings:

- One group prefix mapped to five RPs counts as five group-to-RP mappings.
- Five distinct group prefixes mapped to one RP count as five group-to-RP mappings.

Once the configured limits are reached, no new PIM join messages, PIM register messages, or group-to-RP mappings are accepted unless one of the following occurs:

- You clear the current PIM join states by using the **clear pim join** command. If you use this command on an RP configured for PIM register message limits, the register limit count is also restarted because the PIM join messages are unknown by the RP.



**NOTE:** On the RP, you can also use the **clear pim register** command to clear all of the PIM registers. This command is useful if the current PIM register count is greater than the newly configured PIM register limit. After you clear the PIM registers, new PIM register messages are received up to the configured limit.

- The traffic responsible for the excess PIM join messages and PIM register messages stops and is no longer present.



**CAUTION:** Never restart any of the software processes unless instructed to do so by a customer support engineer.

You restart the PIM routing process on the device. This restart clears all of the configured limits but disrupts routing and therefore requires a maintenance window for the change.

### System Log Messages for PIM Resources

You can optionally configure a system log warning threshold for each of the PIM resources. With this configuration, you can generate and review system log messages to detect if an excessive number of PIM join messages, PIM register messages, or group-to-RP mappings have been received on the device. The system log warning thresholds are configured per PIM resource and are a percentage of the configured maximum limits of the PIM join messages, PIM register messages, and group-to-RP mappings. You can further specify a log interval for each configured PIM resource, which is the amount of time (in seconds) between the log messages.

The log messages convey when the configured limits have been exceeded, when the configured warning thresholds have been exceeded, and when the configured limits drop below the configured warning threshold. [Table 11 on page 548](#) describes the different types of PIM system messages that you might see depending on your system log warning and log interval configurations.

**Table 11: PIM System Log Messages**

System Log Message	Definition
RPD_PIM_SG_THRESHOLD_EXCEED	Records when the (S,G)/(*G) routes exceed the configured warning threshold.
RPD_PIM_REG_THRESH_EXCEED	Records when the PIM registers exceed the configured warning threshold.
RPD_PIM_GRP_RP_MAP_THRES_EXCEED	Records when the group-to-RP mappings exceed the configured warning threshold.
RPD_PIM_SG_LIMIT_EXCEED	Records when the (S,G)/(*G) routes exceed the configured limit, or when the configured log interval has been met and the routes exceed the configured limit.
RPD_PIM_REGISTER_LIMIT_EXCEED	Records when the PIM registers exceed the configured limit, or when the configured log interval has been met and the registers exceed the configured limit.
RPD_PIM_GRP_RP_MAP_LIMIT_EXCEED	Records when the group-to-RP mappings exceed the configured limit, or when the configured log interval has been met and the mapping exceeds the configured limit.
RPD_PIM_SG_LIMIT_BELOW	Records when the (S,G)/(*G) routes drop below the configured limit and the configured log interval.
RPD_PIM_REGISTER_LIMIT_BELOW	Records when the PIM registers drop below the configured limit and the configured log interval.
RPD_PIM_GRP_RP_MAP_LIMIT_BELOW	Records when the group-to-RP mappings drop below the configured limit and the configured log interval.

**See Also**    • [Example: Configuring PIM State Limits on page 548](#)

## Example: Configuring PIM State Limits

This example shows how to set limits on the Protocol Independent Multicast (PIM) state information so that a service provider network can protect itself from potential attacks from misconfigured or misbehaving customer edge (CE) devices and their associated VPN routing and forwarding (VRF) routing instances.

- [Requirements on page 548](#)
- [Overview on page 549](#)
- [Configuration on page 549](#)
- [Verification on page 557](#)

### Requirements

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



## Overview

In this example, a multiprotocol BGP-based multicast VPN (next-generation MBGP MVPN) is configured with limits on the PIM state resources.

The **sglimit maximum** statement sets a limit for the number of accepted (\*G) and (S,G) PIM join states received for the vpn-1 routing instance.

The **rp register-limit maximum** statement configures a limit for the number of PIM register messages received for the vpn-1 routing instance. You configure this statement on the rendezvous point (RP) or on all the devices that might become the RP.


The **group-rp-mapping maximum** statement configures a limit for the number of group-to-RP mappings allowed in the vpn-1 routing instance.

For each configured PIM resource, the **threshold** statement sets a percentage of the maximum limit at which to start generating warning messages in the PIM log file.

For each configured PIM resource, the **log-interval** statement is an amount of time (in seconds) between system log message generation.

Figure 55 on page 549 shows the topology used in this example.

Figure 55: PIM State Limits Topology



“CLI Quick Configuration” on page 549 shows the configuration for all of the devices in Figure 55 on page 549. The section “Step-by-Step Procedure” on page 553 describes the steps on Device PE1.

## 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 CE1**

```

set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30
set interfaces ge-1/2/0 unit 1 family mpls
set interfaces lo0 unit 1 family inet address 192.0.2.1/24
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.1

```

```

set protocols pim rp static address 203.0.113.1
set protocols pim interface all
set routing-options router-id 192.0.2.1

```

Device PE1	<pre> set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30 set interfaces ge-1/2/0 unit 2 family mpls set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30 set interfaces ge-1/2/1 unit 5 family mpls set interfaces vt-1/2/0 unit 2 family inet set interfaces lo0 unit 2 family inet address 192.0.2.2/24 set interfaces lo0 unit 102 family inet address 203.0.113.1/24 set protocols mpls interface ge-1/2/1.5 set protocols bgp group ibgp type internal set protocols bgp group ibgp local-address 192.0.2.2 set protocols bgp group ibgp family inet-vpn any set protocols bgp group ibgp family inet-mvpn signaling set protocols bgp group ibgp neighbor 192.0.2.4 set protocols bgp group ibgp neighbor 192.0.2.5 set protocols ospf area 0.0.0.0 interface lo0.2 passive set protocols ospf area 0.0.0.0 interface ge-1/2/1.5 set protocols ldp interface ge-1/2/1.5 set protocols ldp p2mp set policy-options policy-statement parent_vpn_routes from protocol bgp set policy-options policy-statement parent_vpn_routes then accept set routing-instances vpn-1 instance-type vrf set routing-instances vpn-1 interface ge-1/2/0.2 set routing-instances vpn-1 interface vt-1/2/0.2 set routing-instances vpn-1 interface lo0.102 set routing-instances vpn-1 route-distinguisher 100:100 set routing-instances vpn-1 provider-tunnel ldp-p2mp set routing-instances vpn-1 vrf-target target:1:1 set routing-instances vpn-1 protocols ospf export parent_vpn_routes set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.102 passive set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.2 set routing-instances vpn-1 protocols pim sglimit family inet maximum 100 set routing-instances vpn-1 protocols pim sglimit family inet threshold 70 set routing-instances vpn-1 protocols pim sglimit family inet log-interval 10 set routing-instances vpn-1 protocols pim rp register-limit family inet maximum 100 set routing-instances vpn-1 protocols pim rp register-limit family inet threshold 80 set routing-instances vpn-1 protocols pim rp register-limit family inet log-interval 10 set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet maximum 100 set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet threshold 80 set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet log-interval 10 set routing-instances vpn-1 protocols pim rp static address 203.0.113.1 set routing-instances vpn-1 protocols pim interface ge-1/2/0.2 mode sparse set routing-instances vpn-1 protocols mvpn set routing-options router-id 192.0.2.2 set routing-options autonomous-system 1001 </pre>
Device P	<pre> set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30 set interfaces ge-1/2/0 unit 6 family mpls set interfaces ge-1/2/1 unit 9 family inet address 10.1.1.9/30 set interfaces ge-1/2/1 unit 9 family mpls </pre>

```

set interfaces ge-1/2/2 unit 13 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 192.0.2.3/24
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/1.9
set protocols mpls interface ge-1/2/2.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.9
set protocols ospf area 0.0.0.0 interface ge-1/2/2.13
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/1.9
set protocols ldp interface ge-1/2/2.13
set protocols ldp p2mp
set routing-options router-id 192.0.2.3

```

**Device PE2**

```

set interfaces ge-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 10 family mpls
set interfaces ge-1/2/1 unit 17 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 17 family mpls
set interfaces vt-1/2/0 unit 4 family inet
set interfaces lo0 unit 4 family inet address 192.0.2.4/24
set interfaces lo0 unit 104 family inet address 203.0.113.4/24
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.4
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10
set protocols ldp interface ge-1/2/0.10
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.4
set routing-instances vpn-1 interface ge-1/2/1.17
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.104 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.17
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet maximum 100
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet threshold 80
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet log-interval
  10
set routing-instances vpn-1 protocols pim rp static address 203.0.113.1
set routing-instances vpn-1 protocols pim interface ge-1/2/1.17 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 1001

```

**Device PE3**

```
set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces ge-1/2/1 unit 21 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 21 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 5 family inet address 192.0.2.5/24
set interfaces lo0 unit 105 family inet address 203.0.113.5/24
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
set protocols ldp interface ge-1/2/0.14
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5
set routing-instances vpn-1 interface ge-1/2/1.21
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.105 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.21
set routing-instances vpn-1 protocols pim rp static address 203.0.113.1
set routing-instances vpn-1 protocols pim interface ge-1/2/1.21 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 1001
```

**Device CE2**

```
set interfaces ge-1/2/0 unit 18 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 18 family mpls
set interfaces lo0 unit 6 family inet address 192.0.2.6/24
set protocols sap listen 192.168.0.0
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols pim rp static address 203.0.113.1
set protocols pim interface all
set routing-options router-id 192.0.2.6
```

**Device CE3**

```
set interfaces ge-1/2/0 unit 22 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 22 family mpls
set interfaces lo0 unit 7 family inet address 192.0.2.7/24
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols pim rp static address 203.0.113.1
set protocols pim interface all
set routing-options router-id 192.0.2.7
```

**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 PIM state limits:

1. Configure the network interfaces.

```
[edit interfaces]
user@PE1# set ge-1/2/0 unit 2 family inet address 10.1.1.2/30
user@PE1# set ge-1/2/0 unit 2 family mpls
```

```
user@PE1# set ge-1/2/1 unit 5 family inet address 10.1.1.5/30
user@PE1# set ge-1/2/1 unit 5 family mpls
```

```
user@PE1# set vt-1/2/0 unit 2 family inet
```

```
user@PE1# set lo0 unit 2 family inet address 192.0.2.2/24
user@PE1# set lo0 unit 102 family inet address 203.0.113.1/24
```

2. Configure MPLS on the core-facing interface.

```
[edit protocols mpls]
user@PE1# set interface ge-1/2/1.5
```

3. Configure internal BGP (IBGP) on the main router.

The IBGP neighbors are the other PE devices.

```
[edit protocols bgp group ibgp]
user@PE1# set type internal
user@PE1# set local-address 192.0.2.2
user@PE1# set family inet-vpn any
user@PE1# set family inet-mvpn signaling
user@PE1# set neighbor 192.0.2.4
user@PE1# set neighbor 192.0.2.5
```

4. Configure OSPF on the main router.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface lo0.2 passive
user@PE1# set interface ge-1/2/1.5
```

5. Configure a signaling protocol (RSVP or LDP) on the main router.

```
[edit protocols ldp]
user@PE1# set interface ge-1/2/1.5
user@PE1# set p2mp
```

6. Configure the BGP export policy.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE1# set from protocol bgp
user@PE1# set then accept
```

7. Configure the routing instance.

The customer-facing interfaces and the BGP export policy are referenced in the routing instance.

```
[edit routing-instances vpn-1]
user@PE1# set instance-type vrf

user@PE1# set interface ge-1/2/0.2
user@PE1# set interface vt-1/2/0.2
user@PE1# set interface lo0.102

user@PE1# set route-distinguisher 100:100
user@PE1# set provider-tunnel ldp-p2mp
user@PE1# set vrf-target target:1:1

user@PE1# set protocols ospf export parent_vpn_routes
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.102 passive
user@PE1# set protocols ospf area 0.0.0.0 interface ge-1/2/0.2

user@PE1# set protocols pim rp static address 203.0.113.1
user@PE1# set protocols pim interface ge-1/2/0.2 mode sparse

user@PE1# set protocols mvpn
```

8. Configure the PIM state limits.

```
[edit routing-instances vpn-1 protocols pim]
user@PE1# set sglimit family inet maximum 100
user@PE1# set sglimit family inet threshold 70
user@PE1# set sglimit family inet log-interval 10

user@PE1# set rp register-limit family inet maximum 100
user@PE1# set rp register-limit family inet threshold 80
user@PE1# set rp register-limit family inet log-interval 10

user@PE1# set rp group-rp-mapping family inet maximum 100
user@PE1# set rp group-rp-mapping family inet threshold 80
user@PE1# set rp group-rp-mapping family inet log-interval 10
```

9. Configure the router ID and AS number.

```
[edit routing-options]
user@PE1# set router-id 192.0.2.2
user@PE1# set autonomous-system 1001
```

**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@PE1# show interfaces
ge-1/2/0 {
  unit 2 {
    family inet {
      address 10.1.1.2/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 5 {
    family inet {
      address 10.1.1.5/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 2 {
    family inet;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.0.2.2/24;
    }
  }
  unit 102 {
    family inet {
      address 203.0.113.1/24;
    }
  }
}
user@PE1# show protocols
mpls {
  interface ge-1/2/1.5;
}
bgp {
  group ibgp {
    type internal;
    local-address 192.0.2.2;
    family inet-vpn {
      any;
    }
    family inet-mvpn {
      signaling;
    }
    neighbor 192.0.2.4;
    neighbor 192.0.2.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
```

```
        passive;
    }
    interface ge-1/2/1.5;
}
}
ldp {
    interface ge-1/2/1.5;
    p2mp;
}

user@PE1# show policy-options
policy-statement parent_vpn_routes {
    from protocol bgp;
    then accept;
}

user@PE1# show routing-instances
vpn-1 {
    instance-type vrf;
    interface ge-1/2/0.2;
    interface vt-1/2/0.2;
    interface lo0.102;
    route-distinguisher 100:100;
    provider-tunnel {
        ldp-p2mp;
    }
    vrf-target target:1:1;
    protocols {
        ospf {
            export parent_vpn_routes;
            area 0.0.0.0 {
                interface lo0.102 {
                    passive;
                }
                interface ge-1/2/0.2;
            }
        }
        pim {
            sglimit {
                family inet {
                    maximum 100;
                    threshold 70;
                    log-interval 10;
                }
            }
        }
        rp {
            register-limit {
                family inet {
                    maximum 100;
                    threshold 80;
                    log-interval 10;
                }
            }
            group-rp-mapping {
                family inet {
                    maximum 100;
                    threshold 80;
                }
            }
        }
    }
}
```



```

        log-interval 10;
    }
}
static {
    address 203.0.113.1;
}
}
interface ge-1/2/0.2 {
    mode sparse;
}
}
mvpn;
}
}

user@PE1# show routing-options
router-id 192.0.2.2;
autonomous-system 1001;

```

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

### Verification

Confirm that the configuration is working properly.

#### *Monitoring the PIM State Information*

**Purpose** Verify that the counters are set as expected and are not exceeding the configured limits.

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

```

user@PE1> show pim statistics instance vpn-1
PIM Message type      Received      Sent  Rx errors
V2 Hello                393          390         0
...
V4 (S,G) Maximum                      100
V4 (S,G) Accepted                      0
V4 (S,G) Threshold                     70
V4 (S,G) Log Interval                  10
V4 (grp-prefix, RP) Maximum            100
V4 (grp-prefix, RP) Accepted            0
V4 (grp-prefix, RP) Threshold           80
V4 (grp-prefix, RP) Log Interval        10
V4 Register Maximum                    100
V4 Register Accepted                    0
V4 Register Threshold                   80
V4 Register Log Interval                10

```

**Meaning** The V4 (S,G) Maximum field shows the maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing instance. If this number is met, additional (S,G) entries are not accepted.

The V4 (S,G) Accepted field shows the number of accepted (S,G) IPv4 multicast routes.

The V4 (S,G) Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).

The V4 (S,G) Log Interval field shows the time (in seconds) between consecutive log messages.

The V4 (grp-prefix, RP) Maximum field shows the maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.

The V4 (grp-prefix, RP) Accepted field shows the number of accepted group-to-RP IPv4 multicast mappings.

The V4 (grp-prefix, RP) Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).

The V4 (grp-prefix, RP) Log Interval field shows the time (in seconds) between consecutive log messages.

The V4 Register Maximum field shows the maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.

The V4 Register Accepted field shows the number of accepted IPv4 PIM registers.

The V4 Register Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).

The V4 Register Log Interval field shows the time (in seconds) between consecutive log messages.

**See Also** • [Controlling PIM Resources for Multicast VPNs Overview on page 546](#)

**Related Documentation**

- *Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces*
- *Examples: Configuring the Multicast Forwarding Cache*
- *Example: Configuring MSDP with Active Source Limits and Mesh Groups*

## Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

Selective LSPs are also referred to as selective provider tunnels. Selective provider tunnels carry traffic from some multicast groups in a VPN and extend only to the PE routers that have receivers for these groups. You can configure a selective provider tunnel for group prefixes and source prefixes, or you can use wildcards for the group and source, as described in the Internet draft *draft-rekhter-mvpn-wildcard-spmsi-01.txt*, *Use of Wildcard in S-PMSI Auto-Discovery Routes*.

The following sections describe the scenarios and special considerations when you use wildcards for selective provider tunnels.

- [About S-PMSI on page 559](#)
- [Scenarios for Using Wildcard S-PMSI on page 560](#)
- [Types of Wildcard S-PMSI on page 561](#)
- [Differences Between Wildcard S-PMSI and \(S,G\) S-PMSI on page 561](#)
- [Wildcard \(\\*,\\*\) S-PMSI and PIM Dense Mode on page 561](#)
- [Wildcard \(\\*,\\*\) S-PMSI and PIM-BSR on page 562](#)
- [Wildcard Source and the 0.0.0.0/0 Source Prefix on page 562](#)

### About S-PMSI


The provider multicast service interface (PMSI) is a BGP tunnel attribute that contains the tunnel ID used by the PE router for transmitting traffic through the core of the provider network. A selective PMSI (S-PMSI) autodiscovery route advertises binding of a given MVPN customer multicast flow to a particular provider tunnel. The S-PMSI autodiscovery route advertised by the ingress PE router contains /32 IPv4 or /128 IPv6 addresses for the customer source and the customer group derived from the source-tree customer multicast route.

[Figure 56 on page 560](#) shows a simple MVPN topology. The ingress router, PE1, originates the S-PMSI autodiscovery route. The egress routers, PE2 and PE3, have join state as a result of receiving join messages from CE devices that are not shown in the topology. In response to the S-PMSI autodiscovery route advertisement sent by PE1, PE2, and PE3, elect whether or not to join the tunnel based on the join state. The selective provider tunnel configuration is configured in a VRF instance on PE1.



**NOTE:** The MVPN mode configuration (RPT-SPT or SPT-only) is configured on all three PE routers for all VRFs that make up the VPN. If you omit the MVPN mode configuration, the default mode is SPT-only.

Figure 56: Simple MVPN Topology



### Scenarios for Using Wildcard S-PMSI

A wildcard S-PMSI has the source or the group (or both the source and the group) field set to the wildcard value of 0.0.0.0/0 and advertises binding of multiple customer multicast flows to a single provider tunnel in a single S-PMSI autodiscovery route.

The scenarios under which you might configure a wildcard S-PMSI are as follows:

- When the customer multicast flows are PIM-SM in ASM-mode flows. In this case, a PE router connected to an MVPN customer's site that contains the customer's RP (C-RP) could bind all the customer multicast flows traveling along a customer's RPT tree to a single provider tunnel.
- When a PE router is connected to an MVPN customer's site that contains multiple sources, all sending to the same group.
- When the customer multicast flows are PIM-bidirectional flows. In this case, a PE router could bind to a single provider tunnel all the customer multicast flows for the same group that have been originated within the sites of a given MVPN connected to that PE, and advertise such binding in a single S-PMSI autodiscovery route.
- When the customer multicast flows are PIM-SM in SSM-mode flows. In this case, a PE router could bind to a single provider tunnel all the customer multicast flows coming from a given source located in a site connected to that PE router.
- When you want to carry in the provider tunnel all the customer multicast flows originated within the sites of a given MVPN connected to a given PE router.

## Types of Wildcard S-PMSI

The following types of wildcard S-PMSI are supported:

- A (\*G) S-PMSI matches all customer multicast routes that have the group address. The customer source address in the customer multicast route can be any address, including 0.0.0.0/0 for shared-tree customer multicast routes. A (\*, C-G) S-PMSI autodiscovery route is advertised with the source field set to 0 and the source address length set to 0. The multicast group address for the S-PMSI autodiscovery route is derived from the customer multicast joins.
- A (\*,\*) S-PMSI matches all customer multicast routes. Any customer source address and any customer group address in a customer multicast route can be bound to the (\*,\*) S-PMSI. The S-PMSI autodiscovery route is advertised with the source address and length set to 0 and the group address and length set 0. The remaining fields in the S-PMSI autodiscovery route follow the same rule as (C-S, C-G) S-PMSI, as described in section 12.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt).

## Differences Between Wildcard S-PMSI and (S,G) S-PMSI

For dynamic provider tunnels, each customer multicast stream is bound to a separate provider tunnel, and each tunnel is advertised by a separate S-PMSI autodiscovery route. For static LSPs, multiple customer multicast flows are bound to a single provider tunnel by having multiple S-PMSI autodiscovery routes advertise the same provider tunnel.

When you configure a wildcard (\*G) or (\*,\*) S-PMSI, one or more matching customer multicast routes share a single S-PMSI. All customer multicast routes that have a matching source and group address are bound to the same (\*G) or (\*,\*) S-PMSI and share the same tunnel. The (\*G) or (\*,\*) S-PMSI is established when the first matching remote customer multicast join message is received in the ingress PE router, and deleted when the last remote customer multicast join is withdrawn from the ingress PE router. Sharing a single S-PMSI autodiscovery route improves control plane scalability.

## Wildcard (\*,\*) S-PMSI and PIM Dense Mode

For (S,G) and (\*G) S-PMSI autodiscovery routes in PIM dense mode (PIM-DM), all downstream PE routers receive PIM-DM traffic. If a downstream PE router does not have receivers that are interested in the group address, the PE router instantiates prune state and stops receiving traffic from the tunnel.

Now consider what happens for (\*,\*) S-PMSI autodiscovery routes. If the PIM-DM traffic is not bound by a longer matching (S,G) or (\*G) S-PMSI, it is bound to the (\*,\*) S-PMSI. As is always true for dense mode, PIM-DM traffic is flooded to downstream PE routers over the provider tunnel regardless of the customer multicast join state. Because there is no group information in the (\*,\*) S-PMSI autodiscovery route, egress PE routers join a (\*,\*) S-PMSI tunnel if there is any configuration on the egress PE router indicating interest in PIM-DM traffic.

Interest in PIM-DM traffic is indicated if the egress PE router has one of the following configurations in the VRF instance that corresponds to the instance that imports the S-PMSI autodiscovery route:

- At least one interface is configured in dense mode at the **[edit routing-instances instance-name protocols pim interface]** hierarchy level.
- At least one group is configured as a dense-mode group at the **[edit routing-instances instance-name protocols pim dense-groups group-address]** hierarchy level.

## Wildcard (\*,\*) S-PMSI and PIM-BSR

For (S,G) and (\*,G) S-PMSI autodiscovery routes in PIM bootstrap router (PIM-BSR) mode, an ingress PE router floods the PIM bootstrap message (BSM) packets over the provider tunnel to all egress PE routers. An egress PE router does not join the tunnel unless the message has the ALL-PIM-ROUTERS group. If the message has this group, the egress PE router joins the tunnel, regardless of the join state. The group field in the message determines the presence or absence of the ALL-PIM-ROUTERS address.

Now consider what would happen for (\*,\*) S-PMSI autodiscovery routes used with PIM-BSR mode. If the PIM BSM packets are not bound by a longer matching (S,G) or (\*,G) S-PMSI, they are bound to the (\*,\*) S-PMSI. As is always true for PIM-BSR, BSM packets are flooded to downstream PE routers over the provider tunnel to the ALL-PIM-ROUTERS destination group. Because there is no group information in the (\*,\*) S-PMSI autodiscovery route, egress PE routers always join a (\*,\*) S-PMSI tunnel. Unlike PIM-DM, the egress PE routers might have no configuration suggesting use of PIM-BSR as the RP discovery mechanism in the VRF instance. To prevent all egress PE routers from always joining the (\*,\*) S-PMSI tunnel, the (\*,\*) wildcard group configuration must be ignored.

This means that if you configure PIM-BSR, a wildcard-group S-PMSI can be configured for all other group addresses. The (\*,\*) S-PMSI is not used for PIM-BSR traffic. Either a matching (\*,G) or (S,G) S-PMSI (where the group address is the ALL-PIM-ROUTERS group) or an inclusive provider tunnel is needed to transmit data over the provider core. For PIM-BSR, the longest-match lookup is (S,G), (\*,G), and the inclusive provider tunnel, in that order. If you do not configure an inclusive tunnel for the routing instance, you must configure a (\*,G) or (S,G) selective tunnel. Otherwise, the data is dropped. This is because PIM-BSR functions like PIM-DM, in that traffic is flooded to downstream PE routers over the provider tunnel regardless of the customer multicast join state. However, unlike PIM-DM, the egress PE routers might have no configuration to indicate interest or noninterest in PIM-BSR traffic.

## Wildcard Source and the 0.0.0.0/0 Source Prefix

You can configure a 0.0.0.0/0 source prefix and a wildcard source under the same group prefix in a selective provider tunnel. For example, the configuration might look as follows:

```
routing-instances {
  vpna {
    provider-tunnel {
      selective {
        group 203.0.113.0/24 {
```

```

source 0.0.0.0/0 {
    rsvp-te {
        label-switched-path-template {
            sptnl3;
        }
    }
}
wildcard-source {
    rsvp-te {
        label-switched-path-template {
            sptnl2;
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbps;
}
}
}
}
}
}

```

The functions of the **source 0.0.0.0/0** and **wildcard-source** configuration statements are different. The 0.0.0.0/0 source prefix only matches (C-S, C-G) customer multicast join messages and triggers (C-S, C-G) S-PMSI autodiscovery routes derived from the customer multicast address. Because all (C-S, C-G) join messages are matched by the 0.0.0.0/0 source prefix in the matching group, the wildcard source S-PMSI is used only for (\*C-G) customer multicast join messages. In the absence of a configured 0.0.0.0/0 source prefix, the wildcard source matches (C-S, C-G) and (\*C-G) customer multicast join messages. In the example, a join message for (10.0.1.0/24, 203.0.113.0/24) is bound to **sptnl3**. A join message for (\*, 203.0.113.0/24) is bound to **sptnl2**.

#### Related Documentation

- [Configuring a Selective Provider Tunnel Using Wildcards on page 564](#)
- [Example: Configuring Selective Provider Tunnels Using Wildcards on page 565](#)
- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 540](#)
- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 538](#)

## Configuring a Selective Provider Tunnel Using Wildcards

When you configure a selective provider tunnel for MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), you can use wildcards for the multicast group and source address prefixes. Using wildcards enables a PE router to use a single route to advertise the binding of multiple multicast streams of a given MVPN customer to a single provider's tunnel, as described in

<http://tools.ietf.org/html/draft-rekhter-mvpn-wildcard-spmsi-00>.

Sharing a single route improves control plane scalability because it reduces the number of S-PMSI autodiscovery routes.

To configure a selective provider tunnel using wildcards:

1. Configure a wildcard group matching any group IPv4 address and a wildcard source for (\*,\*) join messages.

```
[edit routing-instances vpna provider-tunnel selective]
user@router# set wildcard-group-inet wildcard-source
```

2. Configure a wildcard group matching any group IPv6 address and a wildcard source for (\*,\*) join messages.

```
[edit routing-instances vpna provider-tunnel selective]
user@router# set wildcard-group-inet6 wildcard-source
```

3. Configure an IP prefix of a multicast group and a wildcard source for (\*,G) join messages.

```
[edit routing-instances vpna provider-tunnel selective]
user@router# set group 203.0.113/24 wildcard-source
```

4. Map the IPv4 join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective wildcard-group-inet
wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective)
label-switched-path-template provider-tunnel1
```

5. Map the IPv6 join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective wildcard-group-inet6
wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective)
label-switched-path-template provider-tunnel2
```

6. Map the (\*,203.0.113/24) join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective group 203.0.113/24
wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective)
label-switched-path-template provider-tunnel3
```



- Related Documentation**
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559](#)
  - [Example: Configuring Selective Provider Tunnels Using Wildcards on page 565](#)
  - [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 540](#)
  - [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 538](#)

## Example: Configuring Selective Provider Tunnels Using Wildcards

With the (\*,G) and (\*,\*) S-PMSI, a customer multicast join message can match more than one S-PMSI. In this case, a customer multicast join message is bound to the longest matching S-PMSI. The longest match is a (S,G) S-PMSI, followed by a (\*,G) S-PMSI and a (\*,\*) S-PMSI, in that order.

Consider the following configuration:

```

routing-instances {
  vpna {
    provider-tunnel {
      selective {
        wildcard-group-inet {
          wildcard-source {
            rsvp-te {
              label-switched-path-template {
                sptnl1;
              }
            }
          }
        }
      }
    }
    group 203.0.113.0/24 {
      wildcard-source {
        rsvp-te {
          label-switched-path-template {
            sptnl2;
          }
        }
      }
    }
    source 10.1.1/24 {
      rsvp-te {
        label-switched-path-template {
          sptnl3;
        }
      }
    }
  }
}

```

For this configuration, the longest-match rule works as follows:

- A customer multicast (10.1.1.1, 203.0.113.1) join message is bound to the sptnl3 S-PMSI autodiscovery route.
- A customer multicast (10.2.1.1, 203.0.113.1) join message is bound to the sptnl2 S-PMSI autodiscovery route.
- A customer multicast (10.1.1.1, 203.1.113.1) join message is bound to the sptnl1 S-PMSI autodiscovery route.

When more than one customer multicast route is bound to the same wildcard S-PMSI, only one S-PMSI autodiscovery route is created. An egress PE router always uses the same matching rules as the ingress PE router that advertises the S-PMSI autodiscovery route. This ensures consistent customer multicast mapping on the ingress and the egress PE routers.

**Related  
Documentation**

- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559](#)
- [Configuring a Selective Provider Tunnel Using Wildcards on page 564](#)

# Providing Resiliency with Redundant Virtual Tunnels

- [Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview on page 567](#)
- [Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 569](#)
- [Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 570](#)

## Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview

In multicast Layer 3 VPNs, virtual tunnel (VT) interfaces are needed to facilitate virtual routing and forwarding (VRF) table lookup based on MPLS labels.

Junos OS supports *redundant VTs* at the Packet Forwarding Engine level to improve resiliency in delivering multicast traffic.



**NOTE:** Redundant VTs are supported only on MX Series routers with MPCs.

To create redundant VTs at the Packet Forwarding Engine level, you add member VTs (vt- interfaces) to a parent VT (rvt interface).

Configuring redundant VT interfaces on MX Series routers with MPCs has the following characteristics:

- When creating a redundant VT, you can add unconfigured VTs as members of the redundant VT. You configure the redundant VT the same way you configure a regular VT, and the member VTs inherit the configuration of the parent redundant VT.
- When a VT with an existing configuration joins a redundant VT, you must configure the redundant VT with the settings from the existing configuration.
- You can create up to 16 redundant VTs.
- You can add up to 32 VTs as members of the redundant VTs.



**NOTE:** The actual number of VTs you can create is determined by the type of chassis and the number of line cards you have.

- When you add more than two VTs to a redundant VT, the members are in active mode by default, and traffic through the redundant VT is load balanced across all members of the redundant VT.
- When you add only two members, you can configure the members in one of two ways:
  - Both members in active mode
  - One member in active mode and the other in backup mode



**BEST PRACTICE:** To set one member to active mode and the other to backup mode, we recommend that the members be hosted on different MPCs. That way, if an MPC fails, it does not bring down the entire rvt interface.

- Class of service and firewall configurations involving VT interfaces work the same on redundant VT interfaces.



**NOTE:** Redundant VTs help to provide resiliency and improve uptime in your network when delivering multicast traffic. When enhanced-ip is enabled at the [edit chassis network-services] hierarchy level, failure of a VT that is a member of a redundant VT can typically be detected and fail over to another member VT provided within 50 milliseconds.

**Related  
Documentation**

- [Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 569](#)
- [Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 570](#)

## Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

In multicast Layer 3 VPNs, virtual tunnel (VT) interfaces are needed to facilitate virtual routing and forwarding (VRF) table lookup based on MPLS labels.

Junos OS supports *redundant VTs* at the Packet Forwarding Engine level to improve resiliency in delivering multicast traffic.



**NOTE:** Redundant VTs are supported only on MX Series routers with MPCs.

To create redundant VTs at the Packet Forwarding Engine level, add member VTs (vt-interfaces) to a redundant VT (rvt interface), and add that redundant VT interface to a vrf routing instance.

To configure a redundant VT:

1. Enable tunnel services on the MX Series router.

```
[edit chassis]
user@host# set fpc slot-number pic number tunnel-services bandwidth (1g | 10g | 20g | 40g)
```

For example:

```
[edit chassis]
user@host# set fpc 3 pic 0 tunnel-services bandwidth 1g
user@host# set fpc 3 pic 2 tunnel-services bandwidth 10g
```

See *Tunnel Interface Configuration on MX Series Routers Overview* for more information.

2. Create the redundant VT interface type.

```
[edit chassis]
user@host# set redundancy-group interface-type redundant-virtual-tunnel
device-count count
```

For example:

```
[edit chassis]
user@host# set redundancy-group interface-type redundant-virtual-tunnel
device-count 4
```

3. Enable **enhanced-ip** under network services.

```
[edit chassis network-services]
user@host# set enhanced-ip
```

4. Create the redundant VT interface.

```
[edit interfaces interface-name]
user@host# set redundancy-group member-interface interface-name
```

For example:

```
[edit interfaces rvt0]
user@host# set redundancy-group member-interface vt-3/0/10
user@host# set redundancy-group member-interface vt-3/2/10
```

See [redundancy-group \(Interfaces\)](#) for more information.



**NOTE:** You configure the remaining parameters for a redundant VT (rvt) interface the same way as you configure a VT (vt-) interface.

5. Add the redundant VT to the appropriate routing instance.

```
[edit routing-instances routing-instance-name]
user@host# set interface interface-name
```

For example:

```
[edit routing-instances pe-vrf]
user@host# set interface rvt0.0
```



**NOTE:** Because the rvt interface contains two or more vt- member interfaces, providing redundancy, adding an rvt interface and a vt- interface to the same routing instance is not supported.

See “[Configuring Routing Instances on PE Routers in VPNs](#)” on page 9 for detailed information about configuring routing instances.

#### Related Documentation

- [Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview on page 567](#)
- [Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 570](#)
- [Tunnel Interface Configuration on MX Series Routers Overview](#)

## Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic

This example shows how to configure redundant virtual tunnels (VTs) in a multiprotocol BGP (MBGP) multicast VPN (MVPN). You configure a virtual loopback tunnel to facilitate virtual routing and forwarding (VRF) table lookup based on MPLS labels. Redundant VTs configured through this method provide near-immediate (less than 50 milliseconds) failover if one of the VTs fails.

- [Requirements on page 571](#)
- [Overview on page 571](#)

- [Configuration on page 572](#)
- [Verification on page 583](#)

## Requirements

This example uses the following hardware and software components:

- MPCs on MX Series routers
- Junos OS Release 15.2

## Overview

When a VT with an existing configuration joins a redundant VT, you must configure the redundant VT with the settings from the existing configuration.

You can add member VTs to a parent VT for redundancy.

On MX Series routers with MPCs, you can configure redundant VTs in these ways:

- You can create up to 16 redundant VTs.
- You can add up to 32 VTs as members of the redundant VTs.



**NOTE:** The actual number of VTs you can create is determined by the type of chassis and the number of line cards you have.

- When you add more than two VTs to a redundant VT, the members are in active mode by default.
- When you add only two members, you can configure the members in one of these ways:
  - Both members in active mode
  - One member in active mode and the other in backup mode


---

## Topology

In this example, Device PE2 has a redundant VT interface configured in a multicast LDP routing instance. The redundant VT interface in this example contains two member VT interfaces, one in active mode and the other in backup mode.

[Figure 57 on page 572](#) shows the topology used in this example.

Figure 57: Redundant VT Interfaces in MBGP MVPN



“CLI Quick Configuration” on page 572 shows the configuration for the customer edge (CE), provider (P), and provider edge (PE) devices in Figure 57 on page 572. See “Step-by-Step Procedure” on page 577 to configure Device PE2.

## 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 CE1**

```

set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set interfaces ge-0/0/0 description "downstream to PE1"
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description "downstream to PE3"
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.17/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description "Source-facing interface"
set interfaces ge-0/0/2 unit 0 family inet address 203.0.113.3/24
set interfaces lo0 unit 0 family inet address 198.51.100.1/24
set protocols bgp group CE-to-PE1 type external
set protocols bgp group CE-to-PE1 export direct
set protocols bgp group CE-to-PE1 peer-as 10
set protocols bgp group CE-to-PE1 local-as 200
set protocols bgp group CE-to-PE1 neighbor 192.0.2.2 family inet unicast
set protocols bgp group CE-to-PE3 type external
set protocols bgp group CE-to-PE3 export direct
set protocols bgp group CE-to-PE3 peer-as 10
set protocols bgp group CE-to-PE3 local-as 200
set protocols bgp group CE-to-PE3 neighbor 192.0.2.18 family inet unicast
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim rp static address 192.168.0.0
set protocols pim interface all

```



```

set protocols pim interface fxp0.0 disable
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-options router-id 198.51.100.1
set routing-options nonstop-routing
set routing-options autonomous-system 200

```

```

Device PE1
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description "to CE1"
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description "to P"
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.5/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces vt-0/0/10 unit 0 family inet
set interfaces vt-0/0/10 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.2/24
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE1-to-PE2 to 198.51.100.4
set protocols mpls label-switched-path PE1-to-PE3 to 198.51.100.6
set protocols mpls label-switched-path p2mp-template-mvpn_1 template
set protocols mpls label-switched-path p2mp-template-mvpn_1 bandwidth 1m
set protocols mpls label-switched-path p2mp-template-mvpn_1 link-protection
set protocols mpls label-switched-path p2mp-template-mvpn_1 p2mp
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface fxp0.0 disable
set protocols bgp local-address 198.51.100.2
set protocols bgp family inet unicast
set protocols bgp family inet-vpn unicast
set protocols bgp family inet-mvpn signaling
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 198.51.100.2
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 198.51.100.4
set protocols bgp group ibgp neighbor 198.51.100.6
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-0/0/0.0
set routing-instances vpn-1 interface vt-0/0/10.0 multicast
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100

```

```

set routing-instances vpn-1 provider-tunnel rsvp-te label-switched-path-template
  p2mp-template-mvpn_1
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 routing-options multipath
set routing-instances vpn-1 routing-options auto-export
set routing-instances vpn-1 protocols bgp export direct
set routing-instances vpn-1 protocols bgp group CE1-PE1 type external
set routing-instances vpn-1 protocols bgp group CE1-PE1 export direct
set routing-instances vpn-1 protocols bgp group CE1-PE1 peer-as 200
set routing-instances vpn-1 protocols bgp group CE1-PE1 neighbor 192.0.2.1 family inet
  unicast
set routing-instances vpn-1 protocols ospf export direct
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface all
set routing-instances vpn-1 protocols pim rp local address 192.168.0.0
set routing-instances vpn-1 protocols pim interface all mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode spt-only
set routing-options router-id 198.51.100.2
set routing-options autonomous-system 10

```

**Device PE3**

```

set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description "to CE1"
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.18/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description "to P"
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.21/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces vt-0/0/10 unit 0 family inet
set interfaces vt-0/0/10 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.6/24
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE3-to-PE1 to 198.51.100.2
set protocols mpls label-switched-path PE3-to-PE2 to 198.51.100.4
set protocols mpls label-switched-path p2mp-template-mvpn_1 template
set protocols mpls label-switched-path p2mp-template-mvpn_1 bandwidth 1m
set protocols mpls label-switched-path p2mp-template-mvpn_1 link-protection
set protocols mpls label-switched-path p2mp-template-mvpn_1 p2mp
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface fxp0.0 disable
set protocols bgp local-address 198.51.100.6
set protocols bgp family inet unicast
set protocols bgp family inet-vpn unicast
set protocols bgp family inet-mvpn signaling
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 198.51.100.6
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 198.51.100.2
set protocols bgp group ibgp neighbor 198.51.100.4
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

```

set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-0/0/0.0
set routing-instances vpn-1 interface vt-0/0/10.0 multicast
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel rsvp-te label-switched-path-template
  p2mp-template-mvpn_1
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 routing-options multipath
set routing-instances vpn-1 routing-options auto-export
set routing-instances vpn-1 protocols bgp export direct
set routing-instances vpn-1 protocols bgp group CE1-PE3 type external
set routing-instances vpn-1 protocols bgp group CE1-PE3 export direct
set routing-instances vpn-1 protocols bgp group CE1-PE3 peer-as 200
set routing-instances vpn-1 protocols bgp group CE1-PE3 neighbor 192.0.2.17 family inet
  unicast
set routing-instances vpn-1 protocols ospf export direct
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface all
set routing-instances vpn-1 protocols pim rp local address 192.168.0.0
set routing-instances vpn-1 protocols pim interface all mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode spt-only
set routing-options router-id 198.51.100.6
set routing-options autonomous-system 10

```

Device P

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description "to PE1"
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.6/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description "to PE2"
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.9/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 description "to PE3"
set interfaces ge-0/0/2 unit 0 family inet address 192.0.2.22/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.3/24
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp local-address 198.51.100.3
set protocols bgp family inet-vpn unicast
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 198.51.100.3
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp cluster 198.51.100.3
set protocols bgp group ibgp neighbor 198.51.100.2
set protocols bgp group ibgp neighbor 198.51.100.4

```

```

set protocols bgp group ibgp neighbor 198.51.100.6
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set routing-options router-id 198.51.100.3
set routing-options autonomous-system 10

```

Device PE2

```

set chassis redundancy-group interface-type redundant-virtual-tunnel device-count 16
set chassis fpc 1 pic 0 tunnel-services bandwidth 1g
set chassis fpc 2 pic 0 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 description "to P"
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.10/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description "to CE2"
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.13/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.4/24
set interfaces rvt0 redundancy-group member-interface vt-1/0/10 active
set interfaces rvt0 redundancy-group member-interface vt-2/0/10 backup
set interfaces rvt0 unit 1000 family inet
set interfaces rvt0 unit 1000 family mpls
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE2-to-PE1 to 198.51.100.2
set protocols mpls label-switched-path PE2-to-PE3 to 198.51.100.6
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 198.51.100.4
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 198.51.100.3
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-0/0/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 interface rvt0.1000
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1

```

```

set routing-instances vpn-1 routing-options multipath
set routing-instances vpn-1 protocols bgp group CE2-PE2 type external
set routing-instances vpn-1 protocols bgp group CE2-PE2 export direct
set routing-instances vpn-1 protocols bgp group CE2-PE2 peer-as 100
set routing-instances vpn-1 protocols bgp group CE2-PE2 neighbor 192.0.2.14 family inet
    unicast
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface all
set routing-instances vpn-1 protocols pim rp static address 192.168.0.0
set routing-instances vpn-1 protocols pim interface all mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode spt-only
set routing-instances vpn-1 protocols mvpn sender-based-rpf
set routing-instances vpn-1 protocols mvpn hot-root-standby source-tree
set routing-options router-id 198.51.100.4
set routing-options autonomous-system 10

```

**Device CE2**

```

set interfaces ge-0/0/0 description "upstream to PE2"
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.14/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 description "Source-facing interface"
set interfaces ge-0/0/1 unit 0 family inet address 203.0.113.2/24
set interfaces lo0 unit 0 family inet address 198.51.100.5/24
set protocols bgp group CE-to-PE type external
set protocols bgp group CE-to-PE export direct
set protocols bgp group CE-to-PE peer-as 10
set protocols bgp group CE-to-PE local-as 100
set protocols bgp group CE-to-PE neighbor 192.0.2.13 family inet unicast
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols pim rp static address 192.168.0.0
set protocols pim interface all
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-options router-id 198.51.100.5
set routing-options nonstop-routing
set routing-options autonomous-system 100

```

**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*.

In this section, the `rvt` interface is configured on Device PE2.

1. Configure the redundant virtual tunnel redundancy group and tunnel services on the chassis.

```

[edit chassis]
user@PE2# set redundancy-group interface-type redundant-virtual-tunnel
    device-count 16
user@PE2# set fpc 1 pic 0 tunnel-services bandwidth 1g
user@PE2# set fpc 2 pic 0 tunnel-services bandwidth 1g
user@PE2# set network-services enhanced-ip

```

2. Configure the physical interfaces and loopback interfaces.

```
[edit interfaces]
user@PE2# set ge-0/0/0 description "to P"
user@PE2# set ge-0/0/0 unit 0 family inet address 192.0.2.10/24
user@PE2# set ge-0/0/0 unit 0 family mpls

user@PE2# set ge-0/0/1 description "to CE2"
user@PE2# set ge-0/0/1 unit 0 family inet address 192.0.2.13/24
user@PE2# set ge-0/0/1 unit 0 family mpls

user@PE2# set lo0 unit 0 family inet address 198.51.100.4/24
```

3. Configure the rvt interface.



**NOTE:** In this example, one member of the rvt interface is set to be active and the other interface is set to be the backup. This approach requires member interfaces to be on separate MPCs.

```
[edit interfaces]
user@PE2# set rvt0 redundancy-group member-interface vt-1/0/10 active
user@PE2# set rvt0 redundancy-group member-interface vt-2/0/10 backup
user@PE2# set rvt0 unit 1000 family inet
user@PE2# set rvt0 unit 1000 family mpls
```

4. Configure MPLS.

```
[edit protocols mpls]
user@PE2# set label-switched-path PE2-to-PE1 to 198.51.100.2
user@PE2# set label-switched-path PE2-to-PE3 to 198.51.100.6
user@PE2# set interface ge-0/0/0.0
user@PE2# set interface fxp0.0 disable
```

5. Configure BGP.

```
[edit protocols bgp]
user@PE2# set group ibgp type internal
user@PE2# set group ibgp local-address 198.51.100.4
user@PE2# set group ibgp family inet any
user@PE2# set group ibgp family inet-vpn any
user@PE2# set group ibgp family inet-mvpn signaling
user@PE2# set group ibgp neighbor 198.51.100.3
```

6. Configure an interior gateway protocol.

```
[edit protocols ospf]
user@PE2# set traffic-engineering
user@PE2# set area 0.0.0.0 interface lo0.0 passive
user@PE2# set area 0.0.0.0 interface ge-0/0/0.0
user@PE2# set area 0.0.0.0 interface fxp0.0 disable
```

7. Configure LDP, PIM, and RSVP.

```
[edit protocols]
user@PE2# set ldp interface all
user@PE2# set ldp interface fxp0.0 disable
user@PE2# set pim interface all
user@PE2# set pim interface fxp0.0 disable
user@PE2# set rsvp interface ge-0/0/0.0
user@PE2# set rsvp interface fxp0.0 disable
```

8. Configure the routing policy.

```
[edit policy-options policy-statement direct]
user@PE2# set from protocol direct
user@PE2# set then accept
```

9. Configure the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set instance-type vrf
user@PE2# set interface ge-0/0/1.0
user@PE2# set interface lo0.1
user@PE2# set route-distinguisher 100:100
user@PE2# set vrf-target target:1:1
user@PE2# set routing-options multipath
user@PE2# set protocols bgp group CE2-PE2 type external
user@PE2# set protocols bgp group CE2-PE2 export direct
user@PE2# set protocols bgp group CE2-PE2 peer-as 100
user@PE2# set protocols bgp group CE2-PE2 neighbor 192.0.2.14 family inet unicast
user@PE2# set protocols ospf area 0.0.0.0 interface all
user@PE2# set protocols pim rp static address 192.168.0.0
user@PE2# set protocols pim interface all mode sparse
user@PE2# set protocols mvpn mvpn-mode spt-only
user@PE2# set protocols mvpn sender-based-rpf
user@PE2# set protocols mvpn hot-root-standby source-tree
```

10. Add the rvt interface to the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set interface rvt0.1000
```

11. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@PE2# set router-id 198.51.100.4
user@PE2# set autonomous-system 10
```

## Results

From configuration mode, confirm your configuration by entering the **show chassis**, **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@PE2# show chassis
redundancy-group {
    interface-type {
        redundant-virtual-tunnel {
            device-count 16;
        }
    }
}
fpc 1 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
        }
    }
}
fpc 2 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
        }
    }
}
network-services {
    enhanced-ip;
}

user@PE2# show interfaces
ge-0/0/0 {
    description "to P";
    unit 0 {
        family inet {
            address 192.0.2.10/24;
        }
        family mpls;
    }
}
ge-0/0/1 {
    description "to CE2";
    unit 0 {
        family inet {
            address 192.0.2.13/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 198.51.100.4/24;
        }
    }
}
rvt0 {
    redundancy-group {
        member-interface vt-1/0/10 {
            active;
        }
        member-interface vt-2/0/10 {

```



```

        backup;
    }
}
unit 1000 {
    family inet;
    family mpls;
}
}
user@PE2# show protocols
mpls {
    label-switched-path PE2-to-PE1 {
        to 198.51.100.2;
    }
    label-switched-path PE2-to-PE3 {
        to 198.51.100.6;
    }
    interface fxp0.0 {
        disable;
    }
    interface ge-0/0/0.0;
}
bgp {
    group ibgp {
        type internal;
        local-address 198.51.100.4;
        family inet {
            any;
        }
        family inet-vpn {
            any;
        }
        family inet-mvpn {
            signaling;
        }
        neighbor 198.51.100.3;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
        interface ge-0/0/0.0;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
pim {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}

```

```
rsvp {
    interface fxp0.0 {
        disable;
    }
    interface ge-0/0/0.0;
}

user@PE2# show policy-options
policy-statement direct {
    from protocol direct;
    then accept;
}

user@PE2# show routing-instances
vpn-1 {
    instance-type vrf;
    interface ge-0/0/1.0;
    interface lo0.1;
    interface rvt0.1000;
    route-distinguisher 100:100;
    vrf-target target:1:1;
    routing-options {
        multipath;
    }
    protocols {
        bgp {
            group CE2-PE2 {
                type external;
                export direct;
                peer-as 100;
                neighbor 192.0.2.14 {
                    family inet {
                        unicast;
                    }
                }
            }
        }
        ospf {
            area 0.0.0.0 {
                interface all;
            }
        }
        pim {
            rp {
                static {
                    address 192.168.0.0;
                }
            }
            interface all {
                mode sparse;
            }
        }
        mvpn {
            mvpn-mode {
                spt-only;
            }
            sender-based-rpf;
            hot-root-standby {
                source-tree;
            }
        }
    }
}
```

```

    }
}

user@PE2# show routing-options
router-id 198.51.100.4;
autonomous-system 10;

```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying the Creation of Virtual Tunnel and Redundant Virtual Tunnel Interfaces on page 583](#)
- [Verifying the Inclusion of the rvt Interface in the Multicast Route on page 583](#)
- [Verifying Traffic Through the rvt Interface and Its Member Interfaces on page 584](#)
- [Verifying Immediate Failover to the Backup Virtual Tunnel on page 585](#)

### Verifying the Creation of Virtual Tunnel and Redundant Virtual Tunnel Interfaces

**Purpose** Verify that the rvt interface is created and that the appropriate member vt- interfaces are contained within the rvt interface.

**Action** From operational mode, enter the **show interfaces terse | match rvt0** command.

```

user@PE2# show interfaces terse | match rvt0
user@host# run show interfaces terse | match rvt0
vt-1/0/10.1000      up    up    container--> rvt0.1000
vt-2/0/10.1000      up    up    container--> rvt0.1000
rvt0                up    up
rvt0.1000           up    up    inet

```

**Meaning** The output shows that **rvt0** was created and that it contains **vt-1/0/10** and **vt-2/0/10** as member interfaces.

### Verifying the Inclusion of the rvt Interface in the Multicast Route

**Purpose** Verify that the **vpn1** multicast routing instance is running and that it contains the **rvt0** interface.

**Action** From operational mode, enter the **show multicast route extensive instance vpn1** command.

```

user@PE2# show multicast route extensive instance vpn1
Instance: vpn1 Family: INET

Group: 203.0.113.3
Source: 10.0.0.2/24

```

```

Upstream rpf interface list:
  rvt0.1000 (P)
    Sender Id: Label 299888
Downstream interface list:
  ge-1/0/5.0
Number of outgoing interfaces: 1
Session description: Unknown
Statistics: 460 kbps, 10000 pps, 4387087 packets
RPF Next-hop ID: 941
Next-hop ID: 1048594
Upstream protocol: MVPN
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
Uptime: 00:07:48

```

**Meaning** The output shows that the **vpn1** routing instance is running and contains **rvt0**.

### Verifying Traffic Through the rvt Interface and Its Member Interfaces

**Purpose** Verify that traffic passes through the **rvt0** interface and its member interfaces as expected. For this section, start streams of unicast and multicast traffic from the source to the receiver.



**NOTE:** In this example, all traffic through the **rvt0** interface is expected to pass through the active **vt-1/0/10** interface, and no traffic is expected to pass through the backup **vt-2/0/10** interface.

**Action** From operational mode, enter the **monitor interface rvt0** command.

```

user@PE2# monitor interface rvt0
Interface: rvt0, Enabled, Link is Up
Encapsulation: VPN-Loopback-tunnel, Speed: 2000mbps
Traffic statistics:
  Input bytes:          2788575982 (14720096 bps)          [10976152]
  Output bytes:          0 (0 bps)                          [0]
  Input packets:        60621217 (40000 pps)              [238612]
  Output packets:       0 (0 pps)                           [0]
Error statistics:
  Input errors:          0                                    [0]
  Input drops:           0                                    [0]
  Input framing errors:  0                                    [0]
  Carrier transitions:   0                                    [0]
  Output errors:         0                                    [0]
  Output drops:          0                                    [0]

```

From this output, **rvt0** is enabled and up and traffic is flowing through it.

Now enter the **monitor interface vt-1/0/10** command.

```
user@PE2# monitor interface vt-1/0/10
Interface: vt-1/0/10, Enabled, Link is Up
Encapsulation: VPN-Loopback-tunnel, Speed: 1000mbps
Traffic statistics:
  Input bytes:          2997120202 (14720096 bps)      [3658702]
  Output bytes:         0 (0 bps)                    [0]
  Input packets:        65154787 (40000 pps)          [79537]
  Output packets:       0 (0 pps)                    [0]
Error statistics:
  Input errors:         0                             [0]
  Input drops:          0                             [0]
  Input framing errors: 0                             [0]
  Carrier transitions:  0                             [0]
  Output errors:        0                             [0]
  Output drops:         0                             [0]
```

From this output, all traffic that is passing through the **rvt0** interface is passing through the active **vt-1/0/10** interface, as expected.

Now enter the **monitor interface vt-2/0/10** command.

```
user@PE2# monitor interface vt-2/0/10
Interface: vt-2/0/10, Enabled, Link is Up
Encapsulation: VPN-Loopback-tunnel, Speed: 1000mbps
Traffic statistics:
  Input bytes:          0 (0 bps)                    [0]
  Output bytes:         0 (0 bps)                    [0]
  Input packets:        0 (0 pps)                    [0]
  Output packets:       0 (0 pps)                    [0]
Error statistics:
  Input errors:         0                             [0]
  Input drops:          0                             [0]
  Input framing errors: 0                             [0]
  Carrier transitions:  0                             [0]
  Output errors:        0                             [0]
  Output drops:         0                             [0]
```

From this output, even though **vt-2/0/10** is enabled and up, no traffic is passing through. This is expected, because the interface is set to **backup**.

**Meaning** While everything is running normally, traffic does pass through the redundant virtual interface, **rvt0**, and all of the traffic passes through its active member interface. If neither member interface is set to be **active** or **backup**, the traffic is expected to be load balanced across both interfaces.

### Verifying Immediate Failover to the Backup Virtual Tunnel

**Purpose** Verify that when the MPC containing the active member interface fails or is restarted, all traffic through the **rvt0** interfaces immediately fails over to the backup member interface.



**BEST PRACTICE:** For this task, we recommend you have one window open showing the live statistics of the backup interface, through the **monitor interface vt-2/0/10** command, and another window from which you can restart the MPC containing the active member interface.

**Action** From operational mode, enter the **request chassis fpc slot 1 restart** command and observe the live statistics of the backup member interface.

```
user@PE2# request chassis fpc slot 1 restart
Interface: vt-2/0/10, Enabled, Link is Up
Encapsulation: VPN-Loopback-tunnel, Speed: 1000mbps
Traffic statistics:
  Input bytes:          354964428 (14720248 bps)      [13220676]
  Output bytes:         0 (0 bps)                    [0]
  Input packets:        7716618 (40000 pps)          [287406]
  Output packets:       0 (0 pps)                    [0]
Error statistics:
  Input errors:         0                             [0]
  Input drops:          0                             [0]
  Input framing errors: 0                             [0]
  Carrier transitions:  0                             [0]
  Output errors:        0                             [0]
  Output drops:         0                             [0]
```

**Meaning** The output shows that the traffic through the **rvt0** interface is immediately fully carried by the backup member interface.

**Related Documentation**

- [Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview on page 567](#)
- [Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 569](#)

# Creating Next Generation MVPN VRF Import and Export Policies

- [Configuring VRF Route Targets for Routing Instances for an MBGP MVPN on page 587](#)
- [Limiting Routes to Be Advertised by an MVPN VRF Instance on page 590](#)

## Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

---

By default, the VPN routing and forwarding (VRF) import and export route targets (configured either using VRF import and export policies or using the **vrf-target** statement) are used for importing and exporting routes with the MBGP MVPN network layer reachability information (NLRI).

You can use the **export-target** and **import-target** statements to override the default VRF import and export route targets. Export and import targets can also be specified specifically for sender sites or receiver sites, or can be borrowed from a configured unicast route target. Note that a sender site export route target is always advertised when security association routes are exported.



**NOTE:** When you configure an MBGP MVPN routing instance, you should not configure a target value for an MBGP MVPN specific route target that is identical to a target value for a unicast route target configured in another routing instance.

Specifying route targets in the MBGP MVPN NLRI for sender and receiver sites is useful when there is a mix of sender only, receiver only, and sender and receiver sites. A sender site route target is used for exporting automatic discovery routes by a sender site and for importing automatic discovery routes by a receiver site. A receiver site route target is used for exporting routes by a receiver site and importing routes by a sender site. A sender and receiver site exports and imports routes with both route targets.

A provider edge (PE) router with sites in a specific MBGP MVPN must determine whether a received automatic discovery route is from a sender site or receiver site based on the following:

- If the PE router is configured to be only in a sender site, route targets are imported only from receiver sites. Imported automatic discovery routes must be from a receiver site.

- If the PE router is configured to be only in a receiver site, route targets are imported only from sender sites. Imported automatic discovery routes must be from a sender site.
- If a PE router is configured to be in both sender sites and receiver sites, these guidelines apply:
  - Along with an import route target, you can optionally configure whether the route target is from a receiver or a sender site.
  - If a configuration is not provided, an imported automatic discovery route is treated as belonging to both the sender site set and the receiver site set.

To configure a route target for the MBGP MVPN routing instance, include the **route-target** statement:

```
route-target {  
  export-target {  
    target target-community;  
    unicast;  
  }  
  import-target {  
    target {  
      target-value;  
      receiver target-value;  
      sender target-value;  
    }  
    unicast {  
      receiver;  
      sender;  
    }  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]

The following sections describes how to configure the export target and the import target for an MBGP MVPN:

- [Configuring the Export Target for an MBGP MVPN on page 588](#)
- [Configuring the Import Target for an MBGP MVPN on page 589](#)

## Configuring the Export Target for an MBGP MVPN

To configure an export target, include the **export-target** statement:

```
export-target {  
  target target-community;  
  unicast;  
}
```



You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route target]

Configure the **target** option to specify the export target community. Configure the **unicast** option to use the same target community that has been specified for unicast.

## Configuring the Import Target for an MBGP MVPN

To configure an import target, include the **import-target** statement:

```
import-target {
  target target-value {
    receiver;
    sender;
  }
  unicast {
    receiver;
    sender;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target]

The following sections describe how to configure the import target and unicast parameters:

- [Configuring the Import Target Receiver and Sender for an MBGP MVPN on page 589](#)
- [Configuring the Import Target Unicast Parameters for an MBGP MVPN on page 590](#)

### Configuring the Import Target Receiver and Sender for an MBGP MVPN

To configure the import target community, include the **target** statement and specify the target community. The target community must be in the format **target:x:y**. The **x** value is either an IP address or an AS number followed by an optional **L** to indicate a 4 byte AS number, and **y** is a number (for example, **target:123456L:100**)

```
target target-value {
  receiver;
  sender;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target]

You can specify the target community used when importing either receiver site sets or sender site sets by including one of the following statements:

- **receiver**—Specify the target community used when importing receiver site sets.
- **sender**—Specify the target community used when importing sender site sets.

---

### Configuring the Import Target Unicast Parameters for an MBGP MVPN

To configure a unicast target community as the import target, include the **unicast** statement:

```
unicast {  
    receiver;  
    sender;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target]

You can specify the unicast target community used when importing either receiver site sets or sender site sets by including one of the following statements:

- **receiver**—Specify the unicast target community used when importing receiver site sets.
- **sender**—Specify the unicast target community used when importing sender site sets.

---

## Limiting Routes to Be Advertised by an MVPN VRF Instance

If a hub-and-spoke deployment uses one VPN routing and forwarding (VRF) routing instance for unicast routing and a separate VRF for MVPN routing, you need to limit the PE routers at the hub site to advertise only IPv4 MVPN routes, only IPv6 MVPN routes, or both. This is necessary to prevent the multicast VRF instance from advertising unicast VPN routes to other PE routers.



**NOTE:** This configuration does not prevent the exportation of VPN routes to other VRF instances on the same router if the **auto-export** statement is included in the [edit routing-options] hierarchy.

---

To configure a VRF routing instance with the name **green** to advertise MVPN routes from both the **inet** and **inet6** address families, perform the following steps:

1. Configure the VRF routing instance to advertise IPv4 routes.

```
user@host# set routing-instances green vrf-advertise-selective family inet-mvpn
```

2. Configure the VRF routing instance to advertise IPv6 routes.

```
user@host# set routing-instances green vrf-advertise-selective family inet6-mvpn
```

After the configuration is committed, only the MVPN routes for the specified address families are advertised from the VRF instance to remote PE routers. To remove the restriction on routes being advertised, delete the **vrf-advertise-selective** statement.



**NOTE:** You cannot include the **vrf-advertise-selective** statement and the **no-vrf-advertise** statement in the same VRF configuration. However, if you configure the **vrf-advertise-selective** statement without any of its options, the router has the same behavior as if you configured the **no-vrf-advertise** statement. VPN routes are prevented from being advertised from a VRF routing instance to the remote PE routers.

**Related  
Documentation**

- [family on page 1056](#)
- [inet-mvpn on page 1072](#)
- [inet6-mvpn on page 1074](#)
- [no-vrf-advertise on page 1107](#)
- [vrf-advertise-selective on page 1152](#)



## CHAPTER 29

# Signaling Provider Tunnels in Next Generation MVPNs

- [PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs on page 593](#)
- [Configuring Nonstop Active Routing for BGP Multicast VPN on page 594](#)
- [Configuring PIM Provider Tunnels for an MBGP MVPN on page 596](#)
- [Configuring PIM-SSM GRE Selective Provider Tunnels on page 597](#)
- [Configuring NLRI Parameters for an MBGP MVPN on page 598](#)

### PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs

You can configure PIM sparse mode, PIM dense mode, auto-RP, and bootstrap router (BSR) for MBGP MVPN networks:

- **PIM sparse mode**—Allows a router to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode includes an explicit join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from the receivers to the rendezvous point (RP).
- **PIM dense mode**—Allows a router to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. Packets are forwarded to all interfaces except the incoming interface. Unlike PIM sparse mode, where explicit joins are required for packets to be transmitted downstream, packets are flooded to all routers in the routing instance in PIM dense mode.
- **Auto-RP**—Uses PIM dense mode to propagate control messages and establish RP mapping. You can configure an auto-RP node in one of three different modes: discovery mode, announce mode, and mapping mode.
- **BSR**—Establishes RPs. A selected router in a network acts as a BSR, which selects a unique RP for different group ranges. BSR messages are flooded using a data tunnel between PE routers.

#### **Related Documentation**

- *Example: Allowing MBGP MVPN Remote Sources*
- *Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN*

## Configuring Nonstop Active Routing for BGP Multicast VPN

---

BGP multicast virtual private network (MVPN) is a Layer 3 VPN application that is built on top of various unicast and multicast routing protocols such as Protocol Independent Multicast (PIM), BGP, RSVP, and LDP. Enabling nonstop active routing (NSR) for BGP MVPN requires that NSR support is enabled for all these protocols.

The state maintained by MVPN includes MVPN routes, cmcast, provider-tunnel, and forwarding information. BGP MVPN NSR synchronizes this MVPN state between the master and backup Routing Engines. While some of the state on the backup Routing Engine is locally built based on the configuration, most of it is built based on triggers from other protocols that MVPN interacts with. The triggers from these protocols are in turn the result of state replication performed by these modules. This includes route change notifications by unicast protocols, join and prune triggers from PIM, remote MVPN route notification by BGP, and provider-tunnel related notifications from RSVP and LDP.

Configuring NSR and unified in-service software upgrade (ISSU) support to the BGP MVPN protocol provides features such as various provider tunnel types, different MVPN modes (source tree, shared-tree), and PIM features. As a result, at the ingress PE, replication is turned on for dynamic LSPs. Thus, when NSR is configured, the state for dynamic LSPs is also replicated to the backup Routing Engine. After the state is resolved on the backup Routing Engine, RSVP sends required notifications to MVPN.

To enable BGP MVPN NSR support, the [advertise-from-main-vpn-tables](#) configuration statement needs to be configured at the **[edit protocols bgp]** hierarchy level.

Nonstop active routing configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. When NSR is configured on a dual Routing Engine platform, the PIM control state is replicated on both Routing Engines.

This PIM state information includes:

- Neighbor relationships
- Join and prune information
- RP-set information
- Synchronization between routes and next hops and the forwarding state between the two Routing Engines

Junos OS supports NSR in the following PIM scenarios:

- Dense mode
- Sparse mode
- SSM
- Static RP
- Auto-RP (for IPv4 only)

- Bootstrap router
- Embedded RP on the non-RP router (for IPv6 only)
- BFD support
- Draft Rosen multicast VPNs and BGP multicast VPNs
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies

Before you begin:

- Configure the router interfaces. See the *Interfaces Fundamentals for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library*.
- Configure a multicast group membership protocol (IGMP or MLD). See *Understanding IGMP* and *Understanding MLD*.
- For this feature to work with IPv6, the routing device must be running Junos OS Release 10.4 or later.

To configure nonstop active routing:

1. Because NSR requires you to configure graceful Routing Engine switchover (GRES), to enable GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

```
[edit]
user@host# set chassis redundancy graceful-switchover
```

2. Include the **synchronize** statement at the **[edit system]** hierarchy level so that configuration changes are synchronized on both Routing Engines.

```
[edit system]
user@host# set synchronize
user@host# exit
```

3. Configure PIM settings on the DR with *sparse mode* and *version*, and *static* address pointing to the rendezvous points.

```
[edit protocols pim]
user@host# set rp static address address
user@host# set interface interface-name mode sparse
user@host# set interface interface-name version 2
```

For example, to set sparse mode, version 2 and static address:

```
[edit protocols pim]
user@host# set rp static address 10.210.255.202
user@host# set interface fe-0/1/3.0 mode sparse
user@host# set interface fe-0/1/3.0 version 2
```

4. Configure per-packet load balancing on the DR.

```
[edit policy-options policy-statement policy-name]  
user@host# set then policy-name per-packet
```

For example, to set load-balance policy:

```
[edit policy-options policy-statement load-balance]  
user@host# set then load-balance per-packet
```

5. Apply the load-balance policy on the DR.

```
[edit]  
user@host# set routing-options forwarding-table export load-balance
```

6. Configure nonstop active routing on the DR.

```
[edit]  
user@host# set routing-options nonstop-routing  
user@host# set routing-options router-id address
```

For example, to set nonstop active routing on the designated router with address 10.210.255.201:

```
[edit]  
user@host# set routing-options router-id 10.210.255.201
```

**Related  
Documentation**

- *Configuring Basic PIM Settings*
- *Understanding Nonstop Active Routing for PIM*

---

## Configuring PIM Provider Tunnels for an MBGP MVPN

---

To configure a Protocol Independent Multicast (PIM) sparse mode provider tunnel for a multicast VPN, include the **pim-asm** statement:

```
pim-asm {  
  group-address address;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]**

To complete the PIM sparse mode provider tunnel configuration, you also need to specify the group address using the **group-address** option. The source address for a PIM sparse mode provider tunnel is configured to be the loopback address of the loopback interface in the inet.0 routing table.



## Configuring PIM-SSM GRE Selective Provider Tunnels

This topic describes how to configure a PIM-SSM GRE selective provider tunnel for an MBGP MVPN.

Creating a selective provider tunnel enables you to move high-rate traffic off the inclusive tunnel and deliver the multicast traffic only to receivers that request it. This improves bandwidth utilization.

To configure a PIM-SSM GRE selective provider tunnel for the 203.0.113.1/24 customer multicast group address, the 10.2.2.2/32 customer source address, and a virtual routing instance named **green**:

1. Configure the multicast group address range to be used for creating selective tunnels. The address prefix can be any valid nonreserved IPv4 multicast address range. Whether you configure a range of addresses or a single address, make sure that you configure enough group addresses for all the selective tunnels needed.

```
user@host# set routing-instances green provider-tunnel selective group 203.0.113.1/24
source 10.2.2.2/32 pim-ssm group-range 192.0.2.0/24
```

2. Configure the threshold rate in kilobits per second (Kbps) for triggering the creation of the selective tunnel. If you set the threshold rate to zero Kbps, the selective tunnel is created immediately, and the multicast traffic does not use an inclusive tunnel at all. Optionally, you can leave the threshold rate unconfigured and the result is the same as setting the threshold to zero.

```
user@host# set routing-instances green provider-tunnel selective group 203.0.113.1/24
source 10.2.2.2/32 threshold-rate 0
```

3. Configure the autonomous system number in the global routing options. This is required in MBGP MVPNs.

```
user@host# set routing-options autonomous-system 100
```

When configuring PIM-SSM GRE selective provider tunnels, keep the following in mind:

- Aggregation of multiple customer multicast routes to a single PIM S-PMSI is not supported.
- Provider tunnel multicast group addresses must be IPv4 addresses, even in configurations in which the customer multicast group and source are IPv6 addresses.

### Related Documentation

- [Multicast VPN Terminology on page 473](#)
- [pim-ssm on page 1110](#)
- [group-range on page 1066](#)
- [threshold-rate on page 1143](#)

## Configuring NLRI Parameters for an MBGP MVPN

To enable VPN signaling where multiprotocol BGP carries multicast VPN NLRI for the IPv4 address family, include the **family inet-mvpn** statement:

```
inet-mvpn {  
  signaling {  
    accepted-prefix-limit {  
      maximum number;  
      teardown percentage {  
        idle-timeout (forever | minutes);  
      }  
    }  
    loops number;  
    prefix-limit {  
      maximum number;  
      teardown percentage {  
        idle-timeout (forever | minutes);  
      }  
    }  
  }  
}
```

To enable VPN signaling where multiprotocol BGP carries multicast VPN NLRI for the IPv6 address family, include the **family inet6-mvpn** statement:

```
inet6-mvpn {  
  signaling {  
    accepted-prefix-limit {  
      maximum number;  
      teardown percentage {  
        idle-timeout (forever | minutes);  
      }  
    }  
    loops number;  
    prefix-limit {  
      maximum number;  
      teardown percentage {  
        idle-timeout (forever | minutes);  
      }  
    }  
  }  
}
```

# Associating Interfaces with MGBP VPNs

- [Understanding Redundant Virtual Tunnel Interfaces in MGBP MVPNs on page 599](#)
- [Example: Configuring Redundant Virtual Tunnel Interfaces in MGBP MVPNs on page 600](#)

## Understanding Redundant Virtual Tunnel Interfaces in MGBP MVPNs

---

In multiprotocol BGP (MGBP) multicast VPNs (MVPNs), VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic).

Starting in Junos OS Release 12.3, you can configure up to eight VT interfaces in a routing instance, thus providing Tunnel PIC redundancy inside the same multicast VPN routing instance. When the active VT interface fails, the secondary one takes over, and you can continue managing multicast traffic with no duplication.

Redundant VT interfaces are supported with RSVP point-to-multipoint provider tunnels as well as multicast LDP provider tunnels. This feature also works for extranets.

You can configure one of the VT interfaces to be the primary interface. If a VT interface is configured as the primary, it becomes the next hop that is used for traffic coming in from the core on the label-switched path (LSP) into the routing instance. When a VT interface is configured to be primary and the VT interface is used for both unicast and multicast traffic, only the multicast traffic is affected.

If no VT interface is configured to be the primary or if the primary VT interface is unusable, one of the usable configured VT interfaces is chosen to be the next hop that is used for traffic coming in from the core on the LSP into the routing instance. If the VT interface in use goes down for any reason, another usable configured VT interface in the routing instance is chosen. When the VT interface in use changes, all multicast routes in the instance also switch their reverse-path forwarding (RPF) interface to the new VT interface to allow the traffic to be received.

To realize the full benefit of redundancy, we recommend that when you configure multiple VT interfaces, at least one of the VT interfaces be on a different Tunnel PIC from the other VT interfaces. However, Junos OS does not enforce this.

## Release History Table

Release	Description
12.3	Starting in Junos OS Release 12.3, you can configure up to eight VT interfaces in a routing instance, thus providing Tunnel PIC redundancy inside the same multicast VPN routing instance.

## Related Documentation

- [Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 600](#)

## Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs

This example shows how to configure redundant virtual tunnel (VT) interfaces in multiprotocol BGP (MBGP) multicast VPNs (MVPNs). To configure, include multiple VT interfaces in the routing instance and, optionally, apply the **primary** statement to one of the VT interfaces.

- [Requirements on page 600](#)
- [Overview on page 600](#)
- [Configuration on page 601](#)
- [Verification on page 608](#)

## Requirements

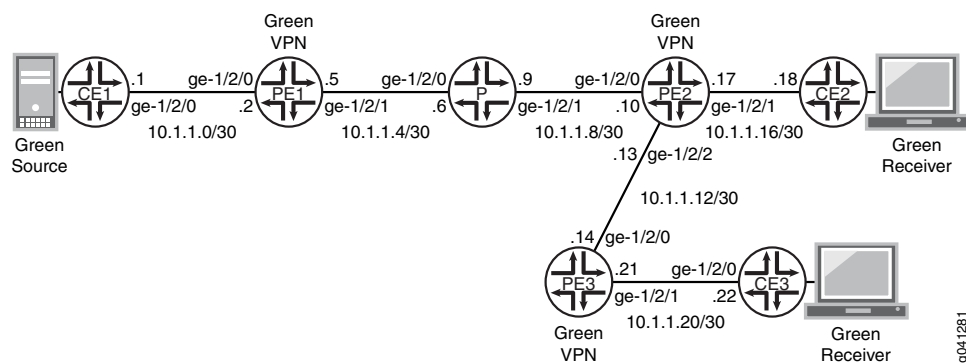
The routing device that has redundant VT interfaces configured must be running Junos OS Release 12.3 or later.

## Overview

In this example, Device PE2 has redundant VT interfaces configured in a multicast LDP routing instance, and one of the VT interfaces is assigned to be the primary interface.

[Figure 57 on page 572](#) shows the topology used in this example.

*Figure 58: Multiple VT Interfaces in MBGP MVPN Topology*



[“CLI Quick Configuration” on page 572](#) shows the configuration for the customer edge (CE), provider (P), and provider edge (PE) devices in [Figure 57 on page 572](#). The section [“Step-by-Step Procedure” on page 604](#) describes the steps on Device PE2.

## Configuration

<b>CLI Quick Configuration</b>	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 <b>[edit]</b> hierarchy level.
<b>Device CE1</b>	<pre> set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.1/30 set interfaces ge-1/2/0 unit 0 family mpls set interfaces lo0 unit 0 family inet address 192.0.2.1/24 set protocols ospf area 0.0.0.0 interface lo0.0 passive set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 set protocols pim rp static address 198.51.100.0 set protocols pim interface all set routing-options router-id 192.0.2.1 </pre>
<b>Device CE2</b>	<pre> set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.18/30 set interfaces ge-1/2/0 unit 0 family mpls set interfaces lo0 unit 0 family inet address 192.0.2.6/24 set protocols sap listen 192.168.0.0 set protocols ospf area 0.0.0.0 interface lo0.0 passive set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 set protocols pim rp static address 198.51.100.0 set protocols pim interface all set routing-options router-id 192.0.2.6 </pre>
<b>Device CE3</b>	<pre> set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.22/30 set interfaces ge-1/2/0 unit 0 family mpls set interfaces lo0 unit 0 family inet address 192.0.2.7/24 set protocols ospf area 0.0.0.0 interface lo0.0 passive set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 set protocols pim rp static address 198.51.100.0 set protocols pim interface all set routing-options router-id 192.0.2.7 </pre>
<b>Device P</b>	<pre> set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.6/30 set interfaces ge-1/2/0 unit 0 family mpls set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.9/30 set interfaces ge-1/2/1 unit 0 family mpls set interfaces lo0 unit 0 family inet address 192.0.2.3/24 set protocols mpls interface ge-1/2/0.0 set protocols mpls interface ge-1/2/1.0 set protocols ospf area 0.0.0.0 interface lo0.0 passive set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 set protocols ospf area 0.0.0.0 interface ge-1/2/1.0 set protocols ldp interface ge-1/2/0.0 set protocols ldp interface ge-1/2/1.0 set protocols ldp p2mp set routing-options router-id 192.0.2.3 </pre>
<b>Device PE1</b>	<pre> set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.2/30 set interfaces ge-1/2/0 unit 0 family mpls </pre>

```

set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.5/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/2/0 unit 2 family inet
set interfaces lo0 unit 0 family inet address 192.0.2.2/24
set interfaces lo0 unit 1 family inet address 198.51.100.0/24
set protocols mpls interface ge-1/2/1.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.2
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ldp interface ge-1/2/1.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-1/2/0.0
set routing-instances vpn-1 interface vt-1/2/0.2 multicast
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set routing-instances vpn-1 protocols pim rp static address 198.51.100.0
set routing-instances vpn-1 protocols pim interface ge-1/2/0.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 1001

```

**Device PE2**

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/1/0 unit 0 family inet
set interfaces vt-1/2/1 unit 0 family inet
set interfaces lo0 unit 0 family inet address 192.0.2.4/24
set interfaces lo0 unit 1 family inet address 203.0.113.4/24
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/2.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.4
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/2.0

```

```

set protocols ldp interface ge-1/2/0.0
set protocols ldp interface ge-1/2/2.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/1/0.0 multicast
set routing-instances vpn-1 interface vt-1/1/0.0 primary
set routing-instances vpn-1 interface vt-1/2/1.0 multicast
set routing-instances vpn-1 interface ge-1/2/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set routing-instances vpn-1 protocols pim rp static address 198.51.100.0
set routing-instances vpn-1 protocols pim interface ge-1/2/1.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 1001

```

**Device PE3**

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 0 family inet address 192.0.2.5/24
set interfaces lo0 unit 1 family inet address 203.0.113.5/24
set protocols mpls interface ge-1/2/0.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5 multicast
set routing-instances vpn-1 interface ge-1/2/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set routing-instances vpn-1 protocols pim rp static address 198.51.100.0
set routing-instances vpn-1 protocols pim interface ge-1/2/1.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.5

```

**set routing-options autonomous-system 1001**

**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 redundant VT interfaces in an MBGP MVPN:

1. Configure the physical interfaces and loopback interfaces.

```
[edit interfaces]
user@PE2# set ge-1/2/0 unit 0 family inet address 10.1.1.10/30
user@PE2# set ge-1/2/0 unit 0 family mpls

user@PE2# set ge-1/2/2 unit 0 family inet address 10.1.1.13/30
user@PE2# set ge-1/2/2 unit 0 family mpls

user@PE2# set ge-1/2/1 unit 0 family inet address 10.1.1.17/30
user@PE2# set ge-1/2/1 unit 0 family mpls

user@PE2# set lo0 unit 0 family inet address 192.0.2.4/24
user@PE2# set lo0 unit 1 family inet address 203.0.113.4/24
```

2. Configure the VT interfaces.

Each VT interface is configurable under one routing instance.

```
[edit interfaces]
user@PE2# set vt-1/1/0 unit 0 family inet
user@PE2# set vt-1/2/1 unit 0 family inet
```

3. Configure MPLS on the physical interfaces.

```
[edit protocols mpls]
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
```

4. Configure BGP.

```
[edit protocols bgp group ibgp]
user@PE2# set type internal
user@PE2# set local-address 192.0.2.4
user@PE2# set family inet-vpn any
user@PE2# set family inet-mvpn signaling
user@PE2# set neighbor 192.0.2.2
user@PE2# set neighbor 192.0.2.5
```

5. Configure an interior gateway protocol.

```
[edit protocols ospf area 0.0.0.0]
user@PE2# set interface lo0.0 passive
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
```



## 6. Configure LDP.

```
[edit protocols ldp]
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
user@PE2# set p2mp
```

## 7. Configure the routing policy.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE2# set from protocol bgp
user@PE2# set then accept
```

## 8. Configure the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set instance-type vrf
user@PE2# set interface ge-1/2/1.0
user@PE2# set interface lo0.1
user@PE2# set route-distinguisher 100:100
user@PE2# set vrf-target target:1:1
user@PE2# set protocols ospf export parent_vpn_routes
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.1 passive
user@PE2# set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
user@PE2# set protocols pim rp static address 198.51.100.0
user@PE2# set protocols pim interface ge-1/2/1.0 mode sparse
user@PE2# set protocols mvpn
```

## 9. Configure redundant VT interfaces in the routing instance.

Make vt-1/1/0.0 the primary interface.

```
[edit routing-instances vpn-1]
user@PE2# set interface vt-1/1/0.0 multicast primary
user@PE2# set interface vt-1/2/1.0 multicast
```

## 10. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@PE2# set router-id 192.0.2.4
user@PE2# set autonomous-system 1001
```

**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@PE2# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      address 10.1.1.10/30;
    }
  }
}
```

```
        family mpls;
    }
}
ge-1/2/2 {
    unit 0 {
        family inet {
            address 10.1.1.13/30;
        }
        family mpls;
    }
}
ge-1/2/1 {
    unit 0 {
        family inet {
            address 10.1.1.17/30;
        }
        family mpls;
    }
}
vt-1/1/0 {
    unit 0 {
        family inet;
    }
}
vt-1/2/1 {
    unit 0 {
        family inet;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.0.2.4/24;
        }
    }
    unit 1 {
        family inet {
            address 203.0.113.4/24;
        }
    }
}

user@PE2# show protocols
mpls {
    interface ge-1/2/0.0;
    interface ge-1/2/2.0;
}
bgp {
    group ibgp {
        type internal;
        local-address 192.0.2.4;
        family inet-vpn {
            any;
        }
        family inet-mvpn {
            signaling;
        }
    }
}
```

```

    }
    neighbor 192.0.2.2;
    neighbor 192.0.2.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface ge-1/2/0.0;
    interface ge-1/2/2.0;
  }
}
ldp {
  interface ge-1/2/0.0;
  interface ge-1/2/2.0;
  p2mp;
}

user@PE2# show policy-options
policy-statement parent_vpn_routes {
  from protocol bgp;
  then accept;
}

user@PE2# show routing-instances
vpn-1 {
  instance-type vrf;
  interface vt-1/1/0.0 {
    multicast;
    primary;
  }
  interface vt-1/2/1.0 {
    multicast;
  }
  interface ge-1/2/1.0;
  interface lo0.1;
  route-distinguisher 100:100;
  vrf-target target:1:1;
  protocols {
    ospf {
      export parent_vpn_routes;
      area 0.0.0.0 {
        interface lo0.1 {
          passive;
        }
        interface ge-1/2/1.0;
      }
    }
    pim {
      rp {
        static {
          address 198.51.100.0;
        }
      }
    }
  }
  interface ge-1/2/1.0 {

```

```

        mode sparse;
    }
}
mvpn;
}
}

user@PE2# show routing-options
router-id 192.0.2.4;
autonomous-system 1001;

```

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

## Verification

Confirm that the configuration is working properly.



**NOTE:** The `show multicast route extensive instance instance-name` command also displays the VT interface in the multicast forwarding table when multicast traffic is transmitted across the VPN.

## Checking the LSP Route

**Purpose** Verify that the expected LT interface is assigned to the LDP-learned route.

**Action** 1. From operational mode, enter the **show route table mpls** command.

```

user@PE2> show route table mpls
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 02:09:36, metric 1
            Receive
1          *[MPLS/0] 02:09:36, metric 1
            Receive
2          *[MPLS/0] 02:09:36, metric 1
            Receive
13         *[MPLS/0] 02:09:36, metric 1
            Receive
299776     *[LDP/9] 02:09:14, metric 1
            > via ge-1/2/0.0, Pop
299776(S=0) *[LDP/9] 02:09:14, metric 1
            > via ge-1/2/0.0, Pop
299792     *[LDP/9] 02:09:09, metric 1
            > via ge-1/2/2.0, Pop
299792(S=0) *[LDP/9] 02:09:09, metric 1
            > via ge-1/2/2.0, Pop
299808     *[LDP/9] 02:09:04, metric 1
            > via ge-1/2/0.0, Swap 299808
299824     *[VPN/170] 02:08:56
            > via ge-1/2/1.0, Pop
299840     *[VPN/170] 02:08:56
            > via ge-1/2/1.0, Pop
299856     *[VPN/170] 02:08:56

```

```

299872          receive table vpn-1.inet.0, Pop
                *[LDP/9] 02:08:54, metric 1
                >   via vt-1/1/0.0, Pop
                via ge-1/2/2.0, Swap 299872

```

2. From configuration mode, change the primary VT interface by removing the **primary** statement from the vt-1/1/0.0 interface and adding it to the vt-1/2/1.0 interface.

```

[edit routing-instances vpn-1]
user@PE2# delete interface vt-1/1/0.0 primary
user@PE2# set interface vt-1/2/1.0 primary
user@PE2# commit

```

3. From operational mode, enter the **show route table mpls** command.

```

user@PE2> show route table mpls
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 02:09:36, metric 1
           Receive
1          *[MPLS/0] 02:09:36, metric 1
           Receive
2          *[MPLS/0] 02:09:36, metric 1
           Receive
13         *[MPLS/0] 02:09:36, metric 1
           Receive
299776     *[LDP/9] 02:09:14, metric 1
           >   via ge-1/2/0.0, Pop
299776(S=0) *[LDP/9] 02:09:14, metric 1
           >   via ge-1/2/0.0, Pop
299792     *[LDP/9] 02:09:09, metric 1
           >   via ge-1/2/2.0, Pop
299792(S=0) *[LDP/9] 02:09:09, metric 1
           >   via ge-1/2/2.0, Pop
299808     *[LDP/9] 02:09:04, metric 1
           >   via ge-1/2/0.0, Swap 299808
299824     *[VPN/170] 02:08:56
           >   via ge-1/2/1.0, Pop
299840     *[VPN/170] 02:08:56
           >   via ge-1/2/1.0, Pop
299856     *[VPN/170] 02:08:56
           receive table vpn-1.inet.0, Pop
299872     *[LDP/9] 02:08:54, metric 1
           >   via vt-1/2/1.0, Pop
           via ge-1/2/2.0, Swap 299872

```

**Meaning** With the original configuration, the output shows the vt-1/1/0.0 interface. If you change the primary interface to vt-1/2/1.0, the output shows the vt-1/2/1.0 interface.

**Related Documentation**

- [Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 599](#)



# Configuring Draft Rosen VPNs

- [Configuring Multicast Layer 3 VPNs on page 611](#)
- [Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN on page 612](#)

## Configuring Multicast Layer 3 VPNs

---

You can configure two types of multicast Layer 3 VPNs using the Junos OS:

- Draft Rosen multicast VPNs—Draft Rosen multicast VPNs are described in RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and based on Section Two of the IETF Internet draft [draft-rosen-vpn-mcast-06.txt](#), *Multicast in MPLS/BGP VPNs* (expired April 2004).
- Next generation multicast VPNs—Next generation multicast VPNs are described in Internet drafts [draft-ietf-l3vpn-2547bis-mcast-bgp-03.txt](#), *BGP Encodings for Multicast in MPLS/BGP IP VPNs* and [draft-ietf-l3vpn-2547bis-mcast-02.txt](#), *Multicast in MPLS/BGP IP VPNs*.

This section describes how to configure draft Rosen multicast VPNs. This information is provided to you in case you already have dual PIM multicast VPNs configured on your network. For information about BGP MPLS multicast VPNs (also known as next generation multicast VPNs), see [“MBGP Multicast VPN Sites” on page 445](#).



**NOTE:** Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

You can configure a Layer 3 VPN to support multicast traffic using the Protocol Independent Multicast (PIM) routing protocol. To support multicast, you need to configure PIM on routers within the VPN and within the service provider's network.

Each PE router configured to run multicast over Layer 3 VPNs must have a Tunnel Services PIC. A Tunnel Services PIC is also required on the P routers that act as rendezvous points (RPs). Tunnel Services PICs are also needed on all the CE routers acting as designated routers (first-hop/last-hop routers) or as RPs, just as they are in non-VPN PIM environments.

Configure the master PIM instance at the **[edit protocols pim]** hierarchy level on the CE and PE routers. This master PIM instance configuration on the PE router should match the configuration on the service providers core routers.

You also need to configure a PIM instance for the Layer 3 VPN at the **[edit routing-instances routing-instance-name protocols pim]** hierarchy level on the PE router. This creates a PIM instance for the indicated routing instance. The configuration of the PIM instance on the PE router should match the PIM instance configured on the CE router the PE router is connected to.

For information about how to configure PIM, see the *Multicast Protocols Feature Guide* .

Include the **vpn-apply-export** statement to configure the group address designated for the VPN in the service provider's network. This address must be unique for each VPN and configured on the VRF routing instance of all PE routers connecting to the same VPN. It ensures that multicast traffic is transmitted only to the specified VPN.

Include the **vpn-apply-export** statement:

**vpn-apply-export** address;

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can include this statement at the following hierarchy levels:

- **[edit routing-instances routing-instance-name protocols pim]**
- **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim]**

The rest of the Layer 3 VPN configuration for multicast is conventional and is described in other sections of this manual. Most of the specific configuration tasks needed to activate multicast in a VPN environment involve PIM.

**Related Documentation**

- *Multicast Protocols Feature Guide*

---

## Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN

---

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics, and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running Draft-Rosen multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across external and internal BGP (EIBGP) upstream paths when the PE router has both external BGP (EBGP) and internal BGP (IBGP) paths toward the source or rendezvous point (RP).

- [Requirements on page 613](#)
- [Overview and Topology on page 613](#)



- [Configuration on page 617](#)
- [Verification on page 620](#)

## Requirements

This example requires the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- Junos OS Release 12.1 or later running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
  - OSPF
  - MPLS
  - LDP
  - PIM
  - BGP
3. Configure a multicast VPN.

## Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across unequal EIBGP routes, if a PE router has EBGp and IBGP paths toward the source (or RP). In previous releases, only the active EBGp path was used to send the join messages. This feature is applicable to IPv4 C-PIM join messages.

During load balancing, if a PE router loses one or more EBGp paths toward the source (or RP), the C-PIM join messages that were previously using the EBGp path are moved to a multicast tunnel interface, and the reverse path forwarding (RPF) neighbor on the multicast tunnel interface is selected based on a hash mechanism.

On discovering the first EBGp path toward the source (or RP), only the new join messages get load-balanced across EIBGP paths, whereas the existing join messages on the multicast tunnel interface remain unaffected.

Though the primary goal for multipath PIM join load balancing is to utilize unequal EIBGP paths for multicast traffic, potential join loops can be avoided if a PE router chooses only the EBGp path when there are one or more join messages for different groups from a remote PE router. If the remote PE router's join message arrives after the PE router has already chosen IBGP as the upstream path, then the potential loops can be broken by changing the selected upstream path to EBGp.



.....

**NOTE:** During a graceful Routing Engine switchover (GRES), the EIBGP path selection for C-PIM join messages can vary, because the upstream interface selection is performed again for the new Routing Engine based on the join messages it receives from the CE and PE neighbors. This can lead to disruption of multicast traffic depending on the number of join messages received and the load on the network at the time of the graceful restart. However, the nonstop active routing feature is not supported and has no impact on the multicast traffic in a Draft-Rosen MVPN scenario.

.....

In this example, PE1 and PE2 are the upstream PE routers for which the multipath PIM join load-balancing feature is configured. Routers PE1 and PE2 have one EBGp path and one IBGP path each toward the source. The Source and Receiver attached to customer edge (CE) routers are Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The existing join-count-based load balancing is performed such that the algorithm first selects the least loaded C-PIM interface. If there is equal or no load on all the C-PIM interfaces, the join messages get distributed equally across the available upstream interfaces.

In [Figure 59 on page 617](#), if the PE1 router receives PIM join messages from the CE2 router, and if there is equal or no load on both the EBGp and IBGP paths toward the source, the join messages get load-balanced on the EIBGP paths.

2. If the selected least loaded interface is a multicast tunnel interface, then there can be a potential join loop if the downstream list of the customer join (C-join) message already contains the multicast tunnel interface. In such a case, the least loaded interface among EBGp paths is selected as the upstream interface for the C-join message.

Assuming that the IBGP path is the least loaded, the PE1 router sends the join messages to PE2 using the IBGP path. If PIM join messages from the PE3 router arrive on PE1, then the downstream list of the C-join messages for PE3 already contains a multicast tunnel interface, which can lead to a potential join loop, because both the upstream and downstream interfaces are multicast tunnel interfaces. In this case, PE1 uses only the EBGp path to send the join messages.

3. If the selected least loaded interface is a multicast tunnel interface and the multicast tunnel interface is not present in the downstream list of the C-join messages, the loop prevention mechanism is not necessary. If any PE router has already advertised data multicast distribution tree (MDT) type, length, and values (TLVs), that PE router is selected as the upstream neighbor.

When the PE1 router sends the join messages to PE2 using the least loaded IBGP path, and if PE3 sends its join messages to PE2, no join loop is created.

4. If no data MDT TLV corresponds to the C-join message, the least loaded neighbor on a multicast tunnel interface is selected as the upstream interface.

On PE routers that have only IBGP paths toward the source (or RP), such as PE3, PIM join load balancing is performed as follows:

1. The PE router only finds a multicast tunnel interface as the RPF interface, and load balancing is done across the C-PIM neighbors on a multicast tunnel interface.

Router PE3 load-balances PIM join messages received from the CE4 router across the IBGP paths to the PE1 and PE2 routers.

2. If any PE router has already advertised data MDT TLVs corresponding to the C-join messages, that PE router is selected as the RPF neighbor.

For a particular C-multicast flow, at least one of the PE routers having EIBGP paths toward the source (or RP) must use only the EBGp path to avoid or break join loops. As a result of the loop avoidance mechanism, a PE router is constrained to choose among EIBGP paths when a multicast tunnel interface is already present in the downstream list.

In [Figure 59 on page 617](#), assuming that the CE2 host is interested in receiving traffic from the Source and CE2 initiates multiple PIM join messages for different groups (Group 1 with group address 203.0.113.1, and Group 2 with group address 203.0.113.2), the join messages for both groups arrive on the PE1 router.

Router PE1 then equally distributes the join messages between the EIBGP paths toward the Source. Assuming that Group 1 join messages are sent to the CE1 router directly using the EBGp path, and Group 2 join messages are sent to the PE2 router using the IBGP path, PE1 and PE2 become the RPF neighbors for Group 1 and Group 2 join messages, respectively.

When the CE3 router initiates Group 1 and Group 2 PIM join messages, the join messages for both groups arrive on the PE2 router. Router PE2 then equally distributes the join messages between the EIBGP paths toward the Source. Since PE2 is the RPF neighbor for Group 2 join messages, it sends the Group 2 join messages directly to the CE1 router using the EBGp path. Group 1 join messages are sent to the PE1 router using the IBGP path.

However, if the CE4 router initiates multiple Group 1 and Group 2 PIM join messages, there is no control over how these join messages received on the PE3 router get distributed to reach the Source. The selection of the RPF neighbor by PE3 can affect PIM join load balancing on EIBGP paths.

- If PE3 sends Group 1 join messages to PE1 and Group 2 join messages to PE2, there is no change in RPF neighbor. As a result, no join loops are created.


- If PE3 sends Group 1 join messages to PE2 and Group 2 join messages to PE1, there is a change in the RPF neighbor for the different groups resulting in the creation of join loops. To avoid potential join loops, PE1 and PE2 do not consider IBGP paths to send the join messages received from the PE3 router. Instead, the join messages are sent directly to the CE1 router using only the EBGp path.

The loop avoidance mechanism in a Draft-Rosen MVPN has the following limitations:

- Because the timing of arrival of join messages on remote PE routers determines the distribution of join messages, the distribution could be sub-optimal in terms of join count.
- Because join loops cannot be avoided and can occur due to the timing of join messages, the subsequent RPF interface change leads to loss of multicast traffic. This can be avoided by implementing the PIM make-before-break feature.

The PIM make-before-break feature is an approach to detect and break C-PIM join loops in a Draft-Rosen MVPN. The C-PIM join messages are sent to the new RPF neighbor after establishing the PIM neighbor relationship, but before updating the related multicast forwarding entry. Though the upstream RPF neighbor would have updated its multicast forwarding entry and started sending the multicast traffic downstream, the downstream router does not forward the multicast traffic (because of RPF check failure) until the multicast forwarding entry is updated with the new RPF neighbor. This helps to ensure that the multicast traffic is available on the new path before switching the RPF interface of the multicast forwarding entry.

Figure 59: PIM Join Load Balancing on Draft-Rosen MVPN



9040919

## 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.

```

PE1 set routing-instances vpn1 instance-type vrf
    set routing-instances vpn1 interface ge-5/0/4.0
    set routing-instances vpn1 interface ge-5/2/0.0
    set routing-instances vpn1 interface lo0.1
    set routing-instances vpn1 route-distinguisher 1:1
    set routing-instances vpn1 vrf-target target:1:1
    set routing-instances vpn1 routing-options multipath vpn-unequal-cost
      equal-external-internal
    set routing-instances vpn1 protocols bgp export direct
    set routing-instances vpn1 protocols bgp group bgp type external
    set routing-instances vpn1 protocols bgp group bgp local-address 192.0.2.4
    set routing-instances vpn1 protocols bgp group bgp family inet unicast
    set routing-instances vpn1 protocols bgp group bgp neighbor 192.0.2.5 peer-as 3
    set routing-instances vpn1 protocols bgp group bgp1 type external
    set routing-instances vpn1 protocols bgp group bgp1 local-address 192.0.2.1
    set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
    set routing-instances vpn1 protocols bgp group bgp1 neighbor 192.0.2.2 peer-as 4
    set routing-instances vpn1 protocols pim vpn-group-address 198.51.100.1
  
```

```

set routing-instances vpn1 protocols pim rp static address 10.255.8.168
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance

```

```

PE2 set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-2/0/3.0
set routing-instances vpn1 interface ge-4/0/5.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 2:2
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
    equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 10.90.10.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.90.10.2 peer-as 45
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 4
set routing-instances vpn1 protocols pim vpn-group-address 198.51.100.1
set routing-instances vpn1 protocols pim rp static address 10.255.8.168
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance

```

**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*. To configure the PE1 router:



**NOTE:** Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing and forwarding (VRF) instance.

```

[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-5/0/4.0
user@PE1# set interface ge-5/2/0.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set vrf-target target:1:1

```

2. Enable protocol-independent load balancing for the VRF instance.

```

[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal

```

3. Configure BGP groups and neighbors to enable PE to CE routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 192.0.2.4
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 192.0.2.5 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 192.0.2.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 192.0.2.2 peer-as 4
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim vpn-group-address 198.51.100.1
user@PE1# set pim rp static address 10.255.8.168
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```

6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

**Results** From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-5/0/4.0;
    interface ge-5/2/0.0;
    interface lo0.1;
    route-distinguisher 1:1;
    vrf-target target:1:1;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
  }
  protocols {
    bgp {
      export direct;
      group bgp {
        type external;
        local-address 192.0.2.4;
        family inet {
          unicast;
        }
      }
    }
  }
}
```

```

        neighbor 192.0.2.5 {
            peer-as 3;
        }
    }
    group bgp1 {
        type external;
        local-address 192.0.2.1;
        family inet {
            unicast;
        }
        neighbor 192.0.2.2 {
            peer-as 4;
        }
    }
}
pim {
    vpn-group-address 198.51.100.1;
    rp {
        static {
            address 10.255.8.168;
        }
    }
    interface all;
    join-load-balance;
}
}
}

```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying PIM Join Load Balancing for Different Groups of Join Messages on page 620](#)

### Verifying PIM Join Load Balancing for Different Groups of Join Messages

**Purpose** Verify PIM join load balancing for the different groups of join messages received on the PE1 router.

**Action** From operational mode, run the **show pim join instance extensive** command.

```

user@PE1>show pim join instance extensive
Instance: PIM.vpn1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 203.0.113.1
Source: *
RP: 10.255.8.168
Flags: sparse,rptree,wildcard
Upstream interface: ge-5/2/0.1
Upstream neighbor: 10.10.10.2
Upstream state: Join to RP

```



```

Downstream neighbors:
  Interface: ge-5/0/4.0
    10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 203.0.113.2
  Source: *
  RP: 10.255.8.168
  Flags: sparse,rptree,wildcard
  Upstream interface: mt-5/0/10.32768
  Upstream neighbor: 19.19.19.19
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: ge-5/0/4.0
      10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 203.0.113.3
  Source: *
  RP: 10.255.8.168
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-5/2/0.1
  Upstream neighbor: 10.10.10.2
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: ge-5/0/4.0
      10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 203.0.113.4
  Source: *
  RP: 10.255.8.168
  Flags: sparse,rptree,wildcard
  Upstream interface: mt-5/0/10.32768
  Upstream neighbor: 19.19.19.19
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: ge-5/0/4.0
      10.40.10.2 State: Join Flags: SRW Timeout: 207

```

**Meaning** The output shows how the PE1 router has load-balanced the C-PIM join messages for four different groups.

- For Group 1 (group address: 203.0.113.1) and Group 3 (group address: 203.0.113.3) join messages, the PE1 router has selected the EBGP path toward the CE1 router to send the join messages.
- For Group 2 (group address: 203.0.113.2) and Group 4 (group address: 203.0.113.4) join messages, the PE1 router has selected the IBGP path toward the PE2 router to send the join messages.

**Related Documentation**

- *PIM Join Load Balancing on Multipath MVPN Routes Overview*
- [Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN on page 623](#)



# Configuring Load Balancing

- [Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN on page 623](#)

## Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN

---

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running next-generation multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across available internal BGP (IBGP) upstream paths when there is no external BGP (EBGP) path present, and across available EBGP upstream paths when external and internal BGP (EIBGP) paths are present toward the source or rendezvous point (RP).

- [Requirements on page 623](#)
- [Overview and Topology on page 624](#)
- [Configuration on page 626](#)
- [Verification on page 630](#)

### Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of M Series, MX Series, or T Series routers.
- Junos OS Release 12.1 running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
  - OSPF
  - MPLS
  - LDP

- PIM
  - BGP
3. Configure a multicast VPN.

## Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across all available IBGP paths when there are only IBGP paths present, and across all available upstream EBGP paths when EIBGP paths are present toward the source (or RP). Unlike Draft-Rosen MVPN, next-generation MVPN does not utilize unequal EIBGP paths to send C-PIM join messages. This feature is applicable to IPv4 C-PIM join messages.

By default, only one active IBGP path is used to send the C-PIM join messages for a PE router having only IBGP paths toward the source (or RP). When there are EIBGP upstream paths present, only one active EBGP path is used to send the join messages.

In a next-generation MVPN, C-PIM join messages are translated into (or encoded as) BGP customer multicast (C-multicast) MVPN routes and advertised with the BGP MCAST-VPN address family toward the sender PE routers. A PE router originates a C-multicast MVPN route in response to receiving a C-PIM join message through its PE router to customer edge (CE) router interface. The two types of C-multicast MVPN routes are:

- Shared tree join route (C-\*, C-G)
  - Originated by receiver PE routers.
  - Originated when a PE router receives a shared tree C-PIM join message through its PE-CE router interface.
- Source tree join route (C-S, C-G)
  - Originated by receiver PE routers.
  - Originated when a PE router receives a source tree C-PIM join message (C-S, C-G), or originated by the PE router that already has a shared tree join route and receives a source active autodiscovery route.

The upstream path in a next-generation MVPN is selected using the Bitwise-XOR hash algorithm as specified in Internet draft draft-ietf-l3vpn-2547bis-mcast, *Multicast in MPLS/BGP IP VPNs*. The hash algorithm is performed as follows:

1. The PE routers in the candidate set are numbered from lower to higher IP address, starting from 0.
2. A bitwise exclusive-or of all the bytes is performed on the C-root (source) and the C-G (group) address.

3. The result is taken modulo  $n$ , where  $n$  is the number of PE routers in the candidate set. The result is **N**.
4. **N** represents the IP address of the upstream PE router as numbered in Step 1.

During load balancing, if a PE router with one or more upstream IBGP paths toward the source (or RP) discovers a new IBGP path toward the same source (or RP), the C-PIM join messages distributed among previously existing IBGP paths get redistributed due to the change in the candidate PE router set.

In this example, PE1, PE2, and PE3 are the PE routers that have the multipath PIM join load-balancing feature configured. Router PE1 has two EBGP paths and one IBGP upstream path, PE2 has one EBGP path and one IBGP upstream path, and PE3 has two IBGP upstream paths toward the Source. Router CE4 is the customer edge (CE) router attached to PE3. Source and Receiver are the Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The C-PIM join messages are sent using EBGP paths only. IBGP paths are not used to propagate the join messages.

In [Figure 60 on page 626](#), the PE1 router distributes the join messages between the two EBGP paths to the CE1 router, and PE2 uses the EBGP path to CE1 to send the join messages.

2. If a PE router loses one or more EBGP paths toward the source (or RP), the RPF neighbor on the multicast tunnel interface is selected based on a hash mechanism.

On discovering the first EBGP path, only new join messages get load-balanced across available EBGP paths, whereas the existing join messages on the multicast tunnel interface are not redistributed.

If the EBGP path from the PE2 router to the CE1 router goes down, PE2 sends the join messages to PE1 using the IBGP path. When the EBGP path to CE1 is restored, only new join messages that arrive on PE2 use the restored EBGP path, whereas join messages already sent on the IBGP path are not redistributed.

On PE routers that have only IBGP paths toward the source (or RP), such as the PE3 router, PIM join load balancing is performed as follows:

1. The C-PIM join messages from CE routers get load-balanced only as BGP C-multicast data messages among IBGP paths.


In [Figure 60 on page 626](#), assuming that the CE4 host is interested in receiving traffic from the Source, and CE4 initiates source join messages for different groups (Group 1 [C-S,C-G1] and Group 2 [C-S,C-G2]), the source join messages arrive on the PE3 router.

Router PE3 then uses the Bitwise-XOR hash algorithm to select the upstream PE router to send the C-multicast data for each group. The algorithm first numbers the upstream PE routers from lower to higher IP address starting from 0.

Assuming that Router PE1 router is numbered **0** and Router PE2 is **1**, and the hash result for Group 1 and Group 2 join messages is **0** and **1**, respectively, the PE3 router selects PE1 as the upstream PE router to send Group 1 join messages, and PE2 as the upstream PE router to send the Group 2 join messages to the Source.

2. The shared join messages for different groups [C-\*,C-G] are also treated in a similar way to reach the destination.

*Figure 60: PIM Join Load Balancing on Next-Generation MVPN*



## 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.

```

PE1 set routing-instances vpn1 instance-type vrf
    set routing-instances vpn1 interface ge-3/0/1.0
    set routing-instances vpn1 interface ge-3/3/2.0
    set routing-instances vpn1 interface lo0.1
    set routing-instances vpn1 route-distinguisher 1:1
    set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
      default-template
    set routing-instances vpn1 vrf-target target:1:1
    set routing-instances vpn1 vrf-table-label
  
```

```

set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.40.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.40.10.2 peer-as 3
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 10.10.10.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.10.10.2 peer-as 3
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bitwise-xor-hash

```

```

PE2 set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-1/0/9.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 2:2
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 3
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bitwise-xor-hash

```

```

PE3 set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/8.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 3:3
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 routing-options autonomous-system 1
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.80.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.80.10.2 peer-as 2
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all

```

```
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bitwise-xor-hash
```

**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*. To configure the PE1 router:



**NOTE:** Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing forwarding (VRF) routing instance.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-3/0/1.0
user@PE1# set interface ge-3/3/2.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set provider-tunnel rsvp-te label-switched-path-template
  default-template
user@PE1# set vrf-target target:1:1
user@PE1# set vrf-table-label
```

2. Enable protocol-independent load balancing for the VRF instance.

```
[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```

3. Configure BGP groups and neighbors to enable PE to CE routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 10.40.10.1
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 10.40.10.2 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 10.10.10.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 10.10.10.2 peer-as 3
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim rp static address 10.255.10.119
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```



6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

7. Configure the mode for C-PIM join messages to use rendezvous-point trees, and switch to the shortest-path tree after the source is known.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-mode rpt-spt
```

8. Configure the VRF instance to use the Bytewise-XOR hash algorithm.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-join-load-balance bytewise-xor-hash
```

## Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show routing-instances
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-3/0/1.0;
    interface ge-3/3/2.0;
    interface lo0.1;
    route-distinguisher 1:1;
    provider-tunnel {
      rsvp-te {
        label-switched-path-template {
          default-template;
        }
      }
    }
    vrf-target target:1:1;
    vrf-table-label;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
    protocols {
      bgp {
        export direct;
        group bgp {
          type external;
          local-address 10.40.10.1;
          family inet {
            unicast;
          }
        }
      }
    }
  }
}
```

```

        neighbor 10.40.10.2 {
            peer-as 3;
        }
    }
    group bgp1 {
        type external;
        local-address 10.10.10.1;
        family inet {
            unicast;
        }
        neighbor 10.10.10.2 {
            peer-as 3;
        }
    }
}
pim {
    rp {
        static {
            address 10.255.10.119;
        }
    }
    interface all;
    join-load-balance;
}
mvpn {
    mvpn-mode {
        rpt-spt;
    }
    mvpn-join-load-balance {
        bitwise-xor-hash;
    }
}
}
}

```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages on page 630](#)

### Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages

**Purpose** Verify MVPN C-multicast route information for different groups of join messages received on the PE3 router.

**Action** From operational mode, run the **show mvpn c-multicast** command.

user@PE3>

```

MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET

Instance : vpn1
MVPN Mode : RPT-SPT
C-mcast IPv4 (S:G)          Ptnl          St
0.0.0.0/0:203.0.113.1/24    RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
192.0.2.2/24:203.0.113.1/24  RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
0.0.0.0/0:203.0.113.2/24    RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
192.0.2.2/24:203.0.113.2/24  RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14

```

**Meaning** The output shows how the PE3 router has load-balanced the C-multicast data for the different groups.

- For source join messages (S,G):
  - 192.0.2.2/24:203.0.113.1/24 (S,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
  - 192.0.2.2/24:203.0.113.2/24 (S,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).
- For shared join messages (\*,G):
  - 0.0.0.0/0:203.0.113.1/24 (\*,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
  - 0.0.0.0/0:203.0.113.2/24 (\*,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).

**Related Documentation** • *PIM Join Load Balancing on Multipath MVPN Routes Overview*



# Tracing Traffic and Operations

- [Tracing MBGP MVPN Traffic and Operations on page 633](#)

## Tracing MBGP MVPN Traffic and Operations

---

To trace MBGP MVPN traffic, you can specify options with the **traceoptions** statement:

1. Specify the name of one or more MVPN trace files using the **file** option for the **traceoptions** at the **[edit routing-instances *routing-instance-name* protocols mvpn]** hierarchy level or at the **[edit protocols mvpn]** hierarchy level:

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

The **file** option includes the following sub-options:

- ***filename***—Specify the name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/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 specified maximum **number** of trace files specified is reached. Then the oldest trace file is overwritten.
  - ***size size***—(Optional) Maximum size of each trace file. 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.
  - ***world-readable | no-world-readable***—(Optional) Enable unrestricted file access or restrict file access to the user who created the file.
2. Specify the **flag** option for the **traceoptions** statement:

```
traceoptions {
  flag flag <flag-modifier> <disable>;
}
```

The following trace flags display the operations associated with multicast VPNs:

- **all**—All multicast VPN tracing options
- **cmcast-join**—Multicast VPN C-multicast join routes

- **error**—Error conditions
- **general**—General events
- **inter-as-ad**—Multicast VPN inter-AS automatic discovery routes
- **intra-as-ad**—Multicast VPN intra-AS automatic discovery routes
- **leaf-ad**—Multicast VPN leaf automatic discovery routes
- **mdt-safi-ad**—Multicast VPN MDT SAFI automatic discovery routes
- **nlri**—Multicast VPN advertisements received or sent by means of BGP
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **source-active**—Multicast VPN source active routes
- **spmsi-ad**—Multicast VPN SPMSI auto discovery active routes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing
- **tunnel**—Provider tunnel events
- **umh**—Upstream multicast hop (UMH) events

**Related  
Documentation**

- [traceoptions on page 1144](#)

## PART 5

# Configuring Full-Mesh, Hub-and-Spoke, and Overlapping VPNs

- [Configuring Full-Mesh VPNs on page 637](#)
- [Configuring Hub-and-Spoke VPNs on page 653](#)
- [Configuring Overlapping VPNs on page 681](#)





# Configuring Full-Mesh VPNs


- [Configuring a Simple Full-Mesh VPN Topology on page 637](#)
- [Configuring a Full-Mesh VPN Topology with Route Reflectors on page 651](#)

## Configuring a Simple Full-Mesh VPN Topology

This example shows how to set up a simple full-mesh service provider VPN configuration, which consists of the following components (see [Figure 61 on page 637](#)):

- Two separate VPNs (VPN-A and VPN-B)
- Two provider edge (PE) routers, both of which service VPN-A and VPN-B
- RSVP as the signaling protocol
- One RSVP label-switched path (LSP) that tunnels between the two PE routers through one provider (P) router

*Figure 61: Example of a Simple VPN Topology*



In this configuration, route distribution in VPN A from Router VPN-A-Paris to Router VPN-A-Tokyo occurs as follows:

1. The customer edge (CE) router VPN-A-Paris announces routes to the PE router Router A.
2. Router A installs the received announced routes into its VPN routing and forwarding (VRF) table, VPN-A.inet.0.
3. Router A creates an MPLS label for the interface between it and Router VPN-A-Paris.
4. Router A checks its VRF export policy.
5. Router A converts the Internet Protocol version 4 (IPv4) routes from Router VPN-A-Paris into VPN IPv4 format using its route distinguisher and announces these routes to PE Router C over the IBGP between the two PE routers.
6. Router C checks its VRF import policy and installs all routes that match the policy into its bgp.l3vpn.0 routing table. (Any routes that do not match are discarded.)
7. Router C checks its VRF import policy and installs all routes that match into its VPN-A.inet.0 routing table. The routes are installed in IPv4 format.
8. Router C announces its routes to the CE router Router VPN-A-Tokyo, which installs them into its master routing table. (For routing platforms running Junos OS, the master routing table is inet.0.)
9. Router C uses the LSP between it and Router A to route all packets from Router VPN-A-Tokyo that are destined for Router VPN-A-Paris.

The final section in this example consolidates the statements needed to configure VPN functionality on each of the service P routers shown in [Figure 61 on page 637](#).



**NOTE:** In this example, a private autonomous system (AS) number is used for the route distinguisher and the route target. This number is used for illustration only. When you are configuring VPNs, you should use an assigned AS number.

---

The following sections explain how to configure the VPN functionality on the PE and P routers. The CE routers have no information about the VPN, so you configure them normally.

- [Enabling an IGP on the PE and P Routers on page 638](#)
- [Enabling RSVP and MPLS on the P Router on page 639](#)
- [Configuring the MPLS LSP Tunnel Between the PE Routers on page 639](#)
- [Configuring IBGP on the PE Routers on page 640](#)
- [Configuring Routing Instances for VPNs on the PE Routers on page 641](#)
- [Configuring VPN Policy on the PE Routers on page 643](#)
- [Simple VPN Configuration Summarized by Router on page 646](#)

## Enabling an IGP on the PE and P Routers

To allow the PE and P routers to exchange routing information among themselves, you must configure an interior gateway protocol (IGP) on all these routers or you must

configure static routes. You configure the IGP on the master instance of the routing protocol process (rpd) (that is, at the **[edit protocols]** hierarchy level), not within the VPN routing instance (that is, not at the **[edit routing-instances]** hierarchy level).

You configure the IGP in the standard way. This configuration example does not include this portion of the configuration.

## Enabling RSVP and MPLS on the P Router

On the P router, Router B, you must configure RSVP and MPLS because this router exists on the MPLS LSP path between the two PE routers, Router A and Router C:

```
[edit]
protocols {
  rsvp {
    interface so-4/0/0.0;
    interface so-6/0/0.0;
  }
  mpls {
    interface so-4/0/0.0;
    interface so-6/0/0.0;
  }
}
```

## Configuring the MPLS LSP Tunnel Between the PE Routers

In this configuration example, RSVP is used for VPN signaling. Therefore, in addition to configuring RSVP, you must enable traffic engineering support in an IGP and you must create an MPLS LSP to tunnel the VPN traffic.

On PE Router A, enable RSVP and configure one end of the MPLS LSP tunnel. In this example, traffic engineering support is enabled for OSPF. When configuring the MPLS LSP, include **interface** statements for all interfaces participating in MPLS, including the interfaces to the PE and CE routers. The statements for the interfaces between the PE and CE routers are needed so that the PE router can create an MPLS label for the private interface. In this example, the first **interface** statement configures MPLS on the interface connected to the LSP, and the remaining three configure MPLS on the interfaces that connect the PE router to the CE routers.

```
[edit]
protocols {
  rsvp {
    interface so-3/0/0.0;
  }
  mpls {
    label-switched-path RouterA-to-RouterC {
      to 10.255.245.47;
    }
    interface so-3/0/0.0;
    interface so-6/0/0.0;
    interface so-6/0/1.0;
    interface ge-0/3/0.0;
  }
  ospf {
```

```

        traffic-engineering;
        area 0.0.0.0 {
            interface so-3/0/0.0;
        }
    }
}

```

On PE Router C, enable RSVP and configure the other end of the MPLS LSP tunnel. Again, traffic engineering support is enabled for OSPF, and you configure MPLS on the interfaces to the LSP and the CE routers.

```

[edit]
protocols {
    rsvp {
        interface so-2/0/0.0;
    }
    mpls {
        label-switched-path RouterC-to-RouterA {
            to 10.255.245.68;
        }
        interface so-2/0/0.0;
        interface ge-1/0/0.0;
        interface at-1/2/0.0;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-2/0/0.0;
        }
    }
}

```

## Configuring IBGP on the PE Routers

On the PE routers, configure an IBGP session with the following properties:

- VPN family—To indicate that the IBGP session is for the VPN, include the **family inet-vpn** statement.
- Loopback address—Include the **local-address** statement, specifying the local PE router's loopback address. The IBGP session for VPNs runs through the loopback address. You must also configure the **lo0** interface at the **[edit interfaces]** hierarchy level. The example does not include this part of the router's configuration.
- Neighbor address—Include the **neighbor** statement, specifying the IP address of the neighboring PE router, which is its loopback (**lo0**) address.

On PE Router A, configure IBGP:

```

[edit]
protocols {
    bgp {
        group PE-RouterA-to-PE-RouterC {
            type internal;
            local-address 10.255.245.68;
            family inet-vpn {

```

```

        unicast;
    }
    neighbor 10.255.245.47;
}
}

```

On PE Router C, configure IBGP:

```

[edit]
protocols {
  bgp {
    group PE-RouterC-to-PE-RouterA {
      type internal;
      local-address 10.255.245.47;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.245.68;
    }
  }
}

```

## Configuring Routing Instances for VPNs on the PE Routers

Both PE routers service VPN-A and VPN-B, so you must configure two routing instances on each router, one for each VPN. For each VPN, you must define the following in the routing instance:

- Route distinguisher, which must be unique for each routing instance on the PE router.
- It is used to distinguish the addresses in one VPN from those in another VPN.
- Instance type of **vrf**, which creates the VRF table on the PE router.
- Interfaces connected to the CE routers.
- VRF import and export policies, which must be the same on each PE router that services the same VPN. Unless an import policy contains only a **then reject** statement, it must include reference to a community. Otherwise, when you try to commit the configuration, the commit fails.



**NOTE:** In this example, a private AS number is used for the route distinguisher. This number is used for illustration only. When you are configuring VPNs, you should use an assigned AS number.

- Routing between the PE and CE routers, which is required for the PE router to distribute VPN-related routes to and from connected CE routers. You can configure a routing protocol—BGP, OSPF, or RIP—or you can configure static routing.

On PE Router A, configure the following routing instance for VPN-A. In this example, Router A uses static routes to distribute routes to and from the two CE routers to which it is connected.

```
[edit]
routing-instance {
  VPN-A-Paris-Munich {
    instance-type vrf;
    interface so-6/0/0.0;
    interface so-6/0/1.0;
    route-distinguisher 65535:0;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    routing-options {
      static {
        route 172.16.0.0/16 next-hop so-6/0/0.0;
        route 172.17.0.0/16 next-hop so-6/0/1.0;
      }
    }
  }
}
```

On PE Router C, configure the following routing instance for VPN-A. In this example, Router C uses BGP to distribute routes to and from the CE router to which it is connected.

```
[edit]
routing-instance {
  VPN-A-Tokyo {
    instance-type vrf;
    interface ge-1/0/0.0;
    route-distinguisher 65535:1;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    protocols {
      bgp {
        group VPN-A-Site2 {
          peer-as 1;
          neighbor 10.12.1.2;
        }
      }
    }
  }
}
```

On PE Router A, configure the following routing instance for VPN-B. In this example, Router A uses OSPF to distribute routes to and from the CE router to which it is connected.

```
[edit]
policy-options {
  policy-statement bgp-to-ospf {
    from {
      protocol bgp;
      route-filter 192.168.1.0/24 orlonger;
    }
    then accept;
  }
}
routing-instance {
  VPN-B-Madrid {
    instance-type vrf;
```

```

interface ge-0/3/0.0;
route-distinguisher 65535:2;
vrf-import VPN-B-import;
vrf-export VPN-B-export;
protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.0 {
      interface ge-0/3/0;
    }
  }
}
}
}

```

On PE Router C, configure the following routing instance for VPN-B. In this example, Router C uses RIP to distribute routes to and from the CE router to which it is connected.

```

[edit]
policy-options {
  policy-statement bgp-to-rip {
    from {
      protocol bgp;
      route-filter 192.168.2.0/24 orlonger;
    }
    then accept;
  }
}
routing-instance {
  VPN-B-Osaka {
    instance-type vrf;
    interface at-1/2/0.0;
    route-distinguisher 65535:3;
    vrf-import VPN-B-import;
    vrf-export VPN-B-export;
    protocols {
      rip {
        group PE-C-to-VPN-B {
          export bgp-to-rip;
          neighbor at-1/2/0;
        }
      }
    }
  }
}
}

```

## Configuring VPN Policy on the PE Routers

Configure the VPN import and export policies on each PE router so that the appropriate routes are installed in the PE router's VRF tables. The VRF table is used to forward packets within a VPN. For VPN-A, the VRF table is VPN-A.inet.0, and for VPN-B it is VPN-B.inet.0.

In the VPN policy, you also configure VPN target communities.

In the following example, a private AS number is used for the route target. This number is used for illustration only. When you are configuring VPNs, you should use an assigned

AS number. The policy qualifiers shown in this example are only those needed for the VPN to function. You can configure additional qualifiers, as needed, for any policies that you configure.

On PE Router A, configure the following VPN import and export policies:

```
[edit]
policy-options {
  policy-statement VPN-A-import {
    term a {
      from {
        protocol bgp;
        community VPN-A;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement VPN-A-export {
    term a {
      from protocol static;
      then {
        community add VPN-A;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  policy-statement VPN-B-import {
    term a {
      from {
        protocol bgp;
        community VPN-B;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement VPN-B-export {
    term a {
      from protocol ospf;
      then {
        community add VPN-B;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
}
```



```

}
community VPN-A members target:65535:4;
community VPN-B members target:65535:5;
}

```

On PE Router C, configure the following VPN import and export policies:

```

[edit]
policy-options {
  policy-statement VPN-A-import {
    term a {
      from {
        protocol bgp;
        community VPN-A;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement VPN-A-export {
    term a {
      from protocol bgp;
      then {
        community add VPN-A;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  policy-statement VPN-B-import {
    term a {
      from {
        protocol bgp;
        community VPN-B;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement VPN-B-export {
    term a {
      from protocol rip;
      then {
        community add VPN-B;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
}

```

```

    }
    community VPN-A members target:65535:4;
    community VPN-B members target:65535:5;
  }

```

To apply the VPN policies on the routers, include the **vrf-export** and **vrf-import** statements when you configure the routing instance. For both VPNs, the VRF import and export policies handle the route distribution across the IBGP session running between the PE routers.

To apply the VPN policies on PE Router A, include the following statements:

```

[edit]
routing-instance {
  VPN-A-Paris-Munich {
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
  }
  VPN-B-Madrid {
    vrf-import VPN-B-import;
    vrf-export VPN-B-export;
  }
}

```

To apply the VPN policies on PE Router C, include the following statements:

```

[edit]
routing-instance {
  VPN-A-Tokyo {
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
  }
  VPN-B-Osaka {
    vrf-import VPN-B-import;
    vrf-export VPN-B-export;
  }
}

```

## Simple VPN Configuration Summarized by Router

### Router A (PE Router)

Routing Instance for VPN-A	<pre> routing-instance {   VPN-A-Paris-Munich {     instance-type vrf;     interface so-6/0/0.0;     interface so-6/0/1.0;     route-distinguisher 65535:0;     vrf-import VPN-A-import;     vrf-export VPN-A-export;   } } </pre>
Instance Routing Protocol	<pre> routing-options {   static {     route 172.16.0.0/16 next-hop so-6/0/0.0;   } } </pre>

	<pre>         route 172.17.0.0/16 next-hop so-6/0/1.0;       }     } </pre>
<b>Routing Instance for VPN-B</b>	<pre> routing-instance {   VPN-B-Madrid {     instance-type vrf;     interface ge-0/3/0.0;     route-distinguisher 65535:2;     vrf-import VPN-B-import;     vrf-export VPN-B-export;   } } </pre>
<b>Instance Routing Protocol</b>	<pre> protocols {   ospf {     area 0.0.0.0 {       interface ge-0/3/0;     }   } } </pre>
<b>Master Protocol Instance</b>	<pre> protocols { } </pre>
<b>Enable RSVP</b>	<pre> rsvp {   interface so-3/0/0.0; } </pre>
<b>Configure an MPLS LSP</b>	<pre> mpls {   label-switched-path RouterA-to-RouterC {     to 10.255.245.47;   }   interface so-3/0/0.0;   interface so-6/0/0.0;   interface so-6/0/1.0;   interface ge-0/3/0.0; } </pre>
<b>Configure IBGP</b>	<pre> bgp {   group PE-RouterA-to-PE-RouterC {     type internal;     local-address 10.255.245.68;     family inet-vpn {       unicast;     }     neighbor 10.255.245.47;   } } </pre>

**Configure OSPF for  
Traffic Engineering  
Support**

```
ospf {  
  traffic-engineering;  
  area 0.0.0.0 {  
    interface so-3/0/0.0;  
  }  
}
```

**Configure VPN Policy**

```
policy-options {  
  policy-statement VPN-A-import {  
    term a {  
      from {  
        protocol bgp;  
        community VPN-A;  
      }  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement VPN-A-export {  
    term a {  
      from protocol static;  
      then {  
        community add VPN-A;  
        accept;  
      }  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement VPN-B-import {  
    term a {  
      from {  
        protocol bgp;  
        community VPN-B;  
      }  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement VPN-B-export {  
    term a {  
      from protocol ospf;  
      then {  
        community add VPN-B;  
        accept;  
      }  
    }  
    term b {  
      then reject;  
    }  
  }  
}
```

```

    }
    community VPN-A members target:65535:4;
    community VPN-B members target:65535:5;
  }

```

### Router B (P Router)

<b>Master Protocol Instance</b>	<pre>protocols { }</pre>
<b>Enable RSVP</b>	<pre>rsvp {   interface so-4/0/0.0;   interface so-6/0/0.0; }</pre>
<b>Enable MPLS</b>	<pre>mpls {   interface so-4/0/0.0;   interface so-6/0/0.0; }</pre>

### Router C (PE Router)

<b>Routing Instance for VPN-A</b>	<pre>routing-instance {   VPN-A-Tokyo {     instance-type vrf;     interface ge-1/0/0.0;     route-distinguisher 65535:1;     vrf-import VPN-A-import;     vrf-export VPN-A-export;   } }</pre>
<b>Instance Routing Protocol</b>	<pre>protocols {   bgp {     group VPN-A-Site2 {       peer-as 1;       neighbor 10.12.1.2;     }   } }</pre>
<b>Routing Instance for VPN-B</b>	<pre>VPN-B-Osaka {   instance-type vrf;   interface at-1/2/0.0;   route-distinguisher 65535:3;   vrf-import VPN-B-import;   vrf-export VPN-B-export; }</pre>
<b>Instance Routing Protocol</b>	<pre>protocols {   rip {</pre>

	<pre>        group PE-C-to-VPN-B {             neighbor at-1/2/0;         }     } }</pre>
Master Protocol Instance	<pre>protocols { }</pre>
Enable RSVP	<pre>rsvp {     interface so-2/0/0.0; }</pre>
Configure an MPLS LSP	<pre>mpls {     label-switched-path RouterC-to-RouterA {         to 10.255.245.68;     }     interface so-2/0/0.0;     interface ge-1/0/0.0;     interface at-1/2/0.0; }</pre>
Configure IBGP	<pre>bgp {     group PE-RouterC-to-PE-RouterA {         type internal;         local-address 10.255.245.47;         family inet-vpn {             unicast;         }         neighbor 10.255.245.68;     } }</pre>
Configure OSPF for Traffic Engineering Support	<pre>ospf {     traffic-engineering;     area 0.0.0.0 {         interface so-2/0/0.0;     } }</pre>
Configure VPN Policy	<pre>policy-options {     policy-statement VPN-A-import {         term a {             from {                 protocol bgp;                 community VPN-A;             }             then accept;         }         term b {             then reject;         }     } }</pre>

```

}
policy-statement VPN-A-export {
  term a {
    from protocol bgp;
    then {
      community add VPN-A;
      accept;
    }
  }
  term b {
    then reject;
  }
}
policy-statement VPN-B-import {
  term a {
    from {
      protocol bgp;
      community VPN-B;
    }
    then accept;
  }
  term b {
    then reject;
  }
}
policy-statement VPN-B-export {
  term a {
    from protocol rip;
    then {
      community add VPN-B;
      accept;
    }
  }
  term b {
    then reject;
  }
}
community VPN-A members target:65535:4;
community VPN-B members target:65535:5;
}

```

## Configuring a Full-Mesh VPN Topology with Route Reflectors

This example is a variation of the full-mesh VPN topology example (described in [“Configuring a Simple Full-Mesh VPN Topology” on page 637](#)) in which one of the PE routers is a BGP route reflector. In this variation, Router C in [Figure 61 on page 637](#) is a route reflector. The only change to its configuration is that you need to include the **cluster** statement when configuring the BGP group:

```

[edit]
protocols {
  bgp {
    group PE-RouterC-to-PE-RouterA {
      type internal;

```

```
    local-address 10.255.245.47;
    family inet-vpn {
        unicast;
    }
    neighbor 10.255.245.68;
    cluster 4.3.2.1;
}
}
```

For the complete configuration example of Router C, see [“Configuring a Full-Mesh VPN Topology with Route Reflectors”](#) on page 651.




# Configuring Hub-and-Spoke VPNs

- [Configuring Hub-and-Spoke VPN Topologies: One Interface on page 653](#)
- [Configuring Hub-and-Spoke VPN Topologies: Two Interfaces on page 665](#)

## Configuring Hub-and-Spoke VPN Topologies: One Interface

Use a one-interface configuration to advertise a default route from a hub or hubs.

*Figure 62: Example of a Hub-and-Spoke VPN Topology with One Interface*



[Figure 62 on page 653](#) illustrates a Layer 3 VPN hub-and-spoke application where there is only one interface between the hub CE (CE1) and the hub PE (PE1). This is the recommended way of configuring hub-and-spoke topologies.

In this configuration, a default route is advertised from the hub to the spokes. If more specific spoke CE routes need to be exchanged between spoke CE routers, then two interfaces are needed between the hub CE and hub PE. See [“Configuring Hub-and-Spoke VPN Topologies: Two Interfaces” on page 665](#) for a two-interface example.

In this configuration example, spoke route distribution is as follows:

1. Spoke CE2 advertises its routes to spoke PE2.
2. Spoke PE2 installs routes from CE2 into its VPN routing and forwarding (VRF) table.
3. Spoke PE2 checks its VRF export policy, adds the route target community, and announces the routes to hub PE1.

4. Hub PE1 checks its VRF import policy and installs routes that match the import policy into table `bgp.l3vpn.0`.
5. Hub PE1 installs routes from table `bgp.l3vpn.0` into the hub VRF table.
6. Hub PE1 announces routes from the hub VRF table to the hub CE1.

In this configuration example, default route distribution is as follows:

1. Hub CE1 announces a default route to hub PE1.
2. Hub PE1 installs the default route into the hub VRF table.
3. Hub PE1 checks its VRF export policy, adds the route target community and announces the default route to spoke PE2 and PE3.
4. Spoke PE2 and PE3 check their VRF import policy and install the default route into table `bgp.l3vpn.0`.
5. Spoke PE2 and PE3 install the routes from table `bgp.l3vpn.0` into their spoke VRF tables.
6. Spoke PE2 and PE3 announce the default route from the spoke VRF table to spoke CE2 and CE3.

The following sections describe how to configure a hub-and-spoke topology with one interface based on the topology illustrated in [Figure 62 on page 653](#):

- [Configuring Hub CE1 on page 654](#)
- [Configuring Hub PE1 on page 655](#)
- [Configuring the P Router on page 655](#)
- [Configuring Spoke PE2 on page 656](#)
- [Configuring Spoke PE3 on page 658](#)
- [Configuring Spoke CE2 on page 659](#)
- [Configuring Spoke CE3 on page 659](#)
- [Enabling Egress Features on the Hub PE Router on page 661](#)

## Configuring Hub CE1

Configure hub CE1 as follows:

```
[edit routing-options]
static {
  route 0.0.0.0/0 discard;
}
autonomous-system 100;
[edit protocols]
bgp {
  group hub {
    type external;
    export default;
    peer-as 200;
    neighbor 10.49.4.1;
  }
}
```

```

}
[edit policy-statement]
default {
  term 1 {
    from {
      protocol static;
      route-filter 0.0.0.0/0 exact;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
}

```

## Configuring Hub PE1

Configure hub PE1 as follows:

```

[edit]
routing-instances {
  hub {
    instance-type vrf;
    interface t3-0/0/0 {
      encapsulation frame-relay;
      unit 0 {
        dlci 16;
        family inet {
          address 10.49.4.1/30;
        }
      }
    }
  }
  vrf-target {
    import target:200:100;
    export target:200:101;
  }
  protocols {
    bgp {
      group hub {
        type external;
        peer-as 100;
        as-override;
        neighbor 10.49.4.2;
      }
    }
  }
}
}

```

## Configuring the P Router

Configure the P Router as follows:

```

[edit]
interfaces {
  t3-0/1/1 {

```

```

    unit 0 {
        family inet {
            address 10.49.2.1/30;
        }
        family mpls;
    }
}
t3-0/1/3 {
    unit 0 {
        family inet {
            address 10.49.0.2/30;
        }
        family mpls;
    }
}
t1-0/2/0 {
    unit 0 {
        family inet {
            address 10.49.1.2/30;
        }
        family mpls;
    }
}
}
[edit]
protocols {
    ospf {
        area 0.0.0.0 {
            interface t3-0/1/3.0;
            interface t1-0/2/0.0;
            interface t3-0/1/1.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
ldp {
    interface t3-0/1/1.0;
    interface t3-0/1/3.0;
    interface t1-0/2/0.0;
}
}

```

## Configuring Spoke PE2

Configure spoke PE2 as follows:

```

[edit]
interfaces {
    t3-0/0/0 {
        unit 0 {
            family inet {
                address 10.49.0.1/30;
            }
            family mpls;
        }
    }
}

```

```

    }
    t1-0/1/2 {
        unit 0 {
            family inet {
                address 10.49.3.1/30;
            }
        }
    }
}
[edit protocols]
bgp {
    group ibgp {
        type internal;
        local-address 10.255.14.182;
        peer-as 200;
        neighbor 10.255.14.176 {
            family inet-vpn {
                unicast;
            }
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface t3-0/0/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface t3-0/0/0.0;
}
[edit]
routing-instances {
    spoke {
        instance-type vrf;
        interface t1-0/1/2.0;
        vrf-target {
            import target:200:101;
            export target:200:100;
        }
        protocols {
            bgp {
                group spoke {
                    type external;
                    peer-as 100;
                    as-override;
                    neighbor 10.49.3.2;
                }
            }
        }
    }
}

```

## Configuring Spoke PE3

Configure spoke PE3 as follows:

```
[edit]
interfaces {
  t3-0/0/0 {
    unit 0 {
      family inet {
        address 10.49.6.1/30;
      }
    }
  }
  t3-0/0/1 {
    unit 0 {
      family inet {
        address 10.49.2.2/30;
      }
      family mpls;
    }
  }
}
[edit protocols]
bgp {
  group ibgp {
    type internal;
    local-address 10.255.14.178;
    peer-as 200;
    neighbor 10.255.14.176 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface t3-0/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface t3-0/0/1.0;
}
[edit]
routing-instances {
  spoke {
    instance-type vrf;
    interface t3-0/0/0.0;
    vrf-target {
      import target:200:101;
      export target:200:100;
    }
  }
}
```

```

protocols {
  bgp {
    group spoke {
      type external;
      peer-as 100;
      as-override;
      neighbor 10.49.6.2;
    }
  }
}

```

## Configuring Spoke CE2

Configure spoke CE2 as follows:

```

[edit routing-options]
autonomous-system 100;
{edit protocols}
  bgp {
    group spoke {
      type external;
      export loopback;
      peer-as 200;
      neighbor 10.49.3.1;
    }
  }
}

```

## Configuring Spoke CE3

Configure spoke CE3 as follows:

```

[edit routing-options]
autonomous-system 100;
[edit protocols]
  bgp {
    group spoke {
      type external;
      export loopback;
      peer-as 200;
      neighbor 10.49.6.1;
    }
  }
}

```

In this configuration example, traffic forwarding is as follows between spoke CE2 and hub CE1:

1. Spoke CE2 forwards traffic using the default route learned from spoke PE2 through BGP.

```

0.0.0.0/0          *[BGP/170] 02:24:15, localpref 100
                   AS path: 200 200 I
                   > to 10.49.3.1 via t1-3/0/1.0

```

2. Spoke PE2 performs a route lookup in the spoke VRF table and forwards the traffic to hub PE2 (through the P router—PE2 pushes two labels) using the default route learned through BGP.

```
0.0.0.0/0          *[BGP/170] 01:35:45, localpref 100, from 10.255.14.176
                   AS path: 100 I
                   > via t3-0/0/1.0, Push 100336, Push 100224(top)
```

3. Hub PE1 does a route lookup in the mpls.0 table for the VPN label **100336**.

```
100336            *[VPN/170] 01:37:03
                   > to 10.49.4.2 via t3-0/0/0.0, Pop
```

4. Hub PE1 forwards the traffic out the interface **t3-0/0/0.0** to hub CE1.

In this configuration example, traffic forwarding is as follows between hub CE1 and spoke CE2:

1. Hub CE1 forwards traffic to the hub PE1 using the route learned through BGP.

```
10.49.10.250/32    *[BGP/170] 02:28:46, localpref 100
                   AS path: 200 200 I
                   > to 10.49.4.1 via t3-3/1/0.0
```

2. Hub PE1 does a route lookup in the hub VRF table and forwards the traffic to spoke PE2 (through the P router—PE1 pushes two labels).

```
10.49.10.250/32    *[BGP/170] 01:41:05, localpref 100, from 10.255.14.182
                   AS path: 100 I
                   > via t1-0/1/0.0, Push 100352, Push 100208(top)
```

3. Spoke PE2 does a route lookup in the mpls.0 table for the VPN label **100352**.

```
100352            *[VPN/170] 02:31:39
                   > to 10.49.3.2 via t1-0/1/2.0, Pop
```

4. Spoke PE2 forwards the traffic out the interface **t1-0/1/2.0** to spoke CE2.

In this configuration example, traffic forwarding is as follows between spoke CE2 and spoke CE3:

1. Spoke CE2 forwards traffic using the default route learned from spoke PE2 through BGP.

```
0.0.0.0/0          *[BGP/170] 02:24:15, localpref 100
                   AS path: 200 200 I
                   > to 10.49.3.1 via t1-3/0/1.0
```

2. Spoke PE2 does a route lookup in the spoke VRF table and forwards the traffic to hub PE1 (through the P router—PE2 pushes two labels) using the default route learned through BGP.



```

0.0.0.0/0      *[BGP/170] 01:35:45, localpref 100, from 10.255.14.176
                AS path: 100 I
                > via t3-0/0/1.0, Push 100336, Push 100224(top)

```

3. Hub PE1 does a route lookup in the mpls.0 table for the VPN label **100336**.

```

100336        *[VPN/170] 01:37:03
                > to 10.49.4.2 via t3-0/0/0.0, Pop

```

4. Hub PE1 forwards the traffic out the interface **t3-0/0/0.0** to the hub CE1.  
 5. Hub CE1 forwards the traffic to hub PE1 using the router learned through BGP.

```

10.49.10.253/32  *[BGP/170] 02:40:03, localpref 100
                  AS path: 200 200 I
                  > to 10.49.4.1 via t3-3/1/0.0

```

6. Hub PE1 does a route lookup in the hub VRF table and forwards the traffic to spoke PE3 (through the P router—PE1 pushes two labels).

```

10.49.10.253/32  *[BGP/170] 01:41:05, localpref 100, from 10.255.14.178
                  AS path: 100 I
                  > via t1-0/1/0.0, Push 100128, Push 100192(top)

```

7. Spoke PE3 does a route lookup in the mpls.0 table for VPN label **100128**.

```

100128        *[VPN/170] 02:41:30
                > to 10.49.6.2 via t3-0/0/0.0, Pop

```

8. Spoke PE3 forwards the traffic out the interface **t3-0/0/0.0** to spoke CE3.

If egress features are needed on the hub PE that require an IP forwarding lookup on the hub VRF routing table, see [“Enabling Egress Features on the Hub PE Router” on page 661](#).

## Enabling Egress Features on the Hub PE Router

This example is provided in conjunction with [“Configuring Hub-and-Spoke VPN Topologies: One Interface” on page 653](#). This example also uses the topology illustrated in [Figure 62 on page 653](#).

If egress features are needed on the hub PE that require an IP forwarding lookup on the hub VRF routing table, the configuration detailed in [“Configuring Hub-and-Spoke VPN Topologies: One Interface” on page 653](#) will not work. Applying the **vrf-table-label** statement on the hub routing instance forces traffic from a remote spoke PE to be forwarded to the hub PE and forces an IP lookup to be performed. Because specific spoke routes are in the hub VRF table, traffic will be forwarded to a spoke PE without going through the hub CE.

The hub PE advertises the default route as follows, using VPN label 1028:

```

hub.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
* 0.0.0.0/0 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 10.255.14.176:2

```

```

VPN Label: 1028
Nexthop: Self
Localpref: 100
AS path: 100 I
Communities: target:200:101

```

Incoming traffic is forwarded using VPN label 1028. The mpls.0 table shows that an IP lookup in the table hub.inet.0 is required:

```

1028          *[VPN/0] 00:00:27
              to table hub.inet.0, Pop

```

However, the hub VRF table hub.inet.0 contains specific spoke routes:

```

10.49.10.250/32  *[BGP/170] 00:00:05, localpref 100, from 10.255.14.182
                  AS path: 100 I
                  > via t1-0/1/0.0, Push 100352, Push 100208(top)
10.49.10.253/32  *[BGP/170] 00:00:05, localpref 100, from 10.255.14.178
                  AS path: 100 I
                  > via t1-0/1/0.0, Push 100128, Push 100192(top)

```

Because of this, traffic is forwarded directly to the spoke PEs without going through the hub CE. To prevent this, you must configure a secondary routing instance for downstream traffic in the hub PE1.

### Configuring Hub PE1

Configure hub PE1 as follows:

```

[edit]
routing-instances {
  hub {
    instance-type vrf;
    interface t3-0/0/0.0;
    vrf-target {
      import target:200:100;
      export target:200:101;
    }
    no-vrf-advertise;
    routing-options {
      auto-export;
    }
    protocols {
      bgp {
        group hub {
          type external;
          peer-as 100;
          as-override;
          neighbor 10.49.4.2;
        }
      }
    }
  }
  hub-downstream {
    instance-type vrf;

```

```

vrf-target target:200:101;
vrf-table-label;
routing-options {
    auto-export;
}
}
}

```

When the **no-vrf-advertise** statement is used at the **[edit routing-instances hub]** hierarchy level, no routing table groups or VRF export policies are required. The **no-vrf-advertise** statement configures the hub PE not to advertise VPN routes from the primary routing-instance **hub**. These routes are instead advertised from the secondary routing instance **hub\_downstream**. See *Junos OS Routing Protocols Library* for more information about the **no-vrf-advertise** statement.

The **auto-export** statement at the **[edit routing-instances hub-downstream routing-options]** hierarchy level identifies routes exported from the hub instance to the hub-downstream instance by looking at the route targets defined for each routing instance. See *Junos OS Routing Protocols Library* for more information about using the **auto-export** statement. See *Configuring Overlapping VPNs Using Automatic Route Export* for more examples of export policy.

With this configuration on hub PE, spoke-to-spoke CE traffic goes through the hub CE and permits egress features (such as filtering) to be enabled on the hub PE.

In this configuration example, traffic forwarding is as follows between spoke CE2 and spoke CE3:

1. Spoke CE2 forwards traffic using the default route learned from spoke PE2 through BGP.

```

0.0.0.0/0          *[BGP/170] 02:24:15, localpref 100
                   AS path: 200 200 I
                   > to 10.49.3.1 via t1-3/0/1.0

```

2. Spoke PE2 does a route lookup in the spoke VRF table and forwards the traffic to hub PE1 (through the P router—PE2 pushes two labels) using the default route learned through BGP.

```

spoke.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:00:09, localpref 100, from 10.255.14.176
                   AS path: 100 I
                   > via t3-0/0/0.0, Push 1029, Push 100224(top)

```

3. Hub PE1 does a route lookup in the mpls.0 table for the VPN label 1029.

```

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
1029              *[VPN/0] 00:11:49
                   to table hub_downstream.inet.0, Pop

```

The VPN label **1029** is advertised because:

- a. The **vrf-table-label** statement is applied at the **[edit routing-instances hub\_downstream]** hierarchy level in the hub PE1 configuration.
- b. The **no-vrf-advertise** statement is applied at the **[edit routing-instances hub]** hierarchy level, instructing the router to advertise the route from the secondary table.

Therefore, IP lookups are performed in the **hub\_downstream.inet.0** table, not in the **hub.inet.0** table.

Issue the **show route advertising-protocol** command on the hub PE to a spoke PE to verify the VPN label **1029** advertisement:

```
user@host> show route advertising-protocol
hub_downstream.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
  * 0.0.0.0/0 (1 entry, 1 announced)
    BGP group ibgp type Internal
      Route Distinguisher: 10.255.14.176:3
      VPN Label: 1029
      Nexthop: Self
      Localpref: 100
      AS path: 100 I
      Communities: target:200:101
```

4. Hub PE1 performs an IP lookup in the **hub\_downstream.inet.0** table and forwards the traffic out interface **t3-0/0/0.0** to hub CE1.

```
hub_downstream.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
  0.0.0.0/0 (1 entry, 1 announced)
    *BGP Preference: 170/-101
      Next-hop reference count: 4
      Source: 10.49.4.2
      Next hop: 10.49.4.2 via t3-0/0/0.0, selected
      State: <Secondary Active Ext>
      Peer AS: 100
      Age: 3:03
      Task: BGP_100.10.49.4.2+1707
      Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
      AS path: 100 I
      Communities: target:200:101
      Localpref: 100
      Router ID: 10.49.10.251
      Primary Routing Table hub.inet.0
```

The primary routing table is **hub.inet.0**, indicating that this route was exported from table **hub.inet.0** into this **hub\_downstream.inet.0** table as a result of the **no-vrf-advertise** statement at the **[edit routing-instances hub]** hierarchy level and the **auto-export** statement at the **[edit routing-instances hub-downstream routing-options]** hierarchy level in the hub PE1 configuration.

5. Hub CE1 forwards the traffic back to hub PE1 using the router learned through BGP.

```
10.49.10.253/32    *[BGP/170] 02:40:03, localpref 100
                  AS path: 200 200 I
                  > to 10.49.4.1 via t3-3/1/0.0
```

6. Hub PE1 performs a route lookup in the hub VRF table and forwards the traffic to spoke PE3 (through the P router—PE1 pushes two labels).

```
10.49.10.253/32      *[BGP/170] 01:41:05, localpref 100, from 10.255.14.178
                    AS path: 100 I
                    > via t1-0/1/0.0, Push 100128, Push 100192(top)
```

7. Spoke PE3 performs a route lookup in the mpls.0 table for VPN label **100128**.

```
100128              *[VPN/170] 02:41:30
                    > to 10.49.6.2 via t3-0/0/0.0, Pop
```

8. Spoke PE3 forwards traffic out interface **t3-0/0/0.0** to spoke CE3.


## Configuring Hub-and-Spoke VPN Topologies: Two Interfaces

Use a two-interface configuration to propagate routes from spoke to spoke.

The example in this section configures a hub-and-spoke topology with two interfaces using the following components (see [Figure 63 on page 666](#)):

- One hub PE router (Router D).
- One hub CE router connected to the hub PE router. For this hub-and-spoke VPN topology to function properly, there must be two interfaces connecting the hub PE router to the hub CE router, and each interface must have its own VRF table on the PE router:
  - The first interface (here, interface ge-0/0/0.0) is used to announce spoke routes to the hub CE router. The VRF table associated with this interface contains the routes being announced by the spoke PE routers to the hub CE router.
  - The second interface (here, interface ge-0/0/1.0) is used to receive route announcements from the hub CE that are destined for the hub-and-spoke routers. The VRF table associated with this interface contains the routes announced by the hub CE router to the spoke PE routers. For this example, two separate physical interfaces are used. It would also work if you were to configure two separate logical interfaces sharing the same physical interface between the hub PE router and the hub CE router.
- Two spoke PE routers (Router E and Router F).
- Two spoke CE routers (CE1 and CE2), one connected to each spoke PE router.
- LDP as the signaling protocol.

Figure 63: Example of a Hub-and-Spoke VPN Topology with Two Interfaces




In this configuration, route distribution from spoke CE Router CE1 occurs as follows:

1. Spoke Router CE1 announces its routes to spoke PE Router E.
2. Router E installs the routes from CE1 into its VRF table.
3. After checking its VRF export policy, Router E adds the spoke target community to the routes from Router CE1 that passed the policy and announces them to the hub PE router, Router D.
4. Router D checks the VRF import policy associated with interface ge-0/0/0.0 and places all routes from spoke PE routers that match the policy into its bgp.l3vpn routing table. (Any routes that do not match are discarded.)
5. Router D checks its VRF import policy associated with interface ge-0/0/0.0 and installs all routes that match into its spoke VRF table. The routes are installed with the spoke target community.
6. Router D announces routes to the hub CE over interface ge-0/0/0.
7. The hub CE router announces the routes back to the hub PE Router D over the second interface to the hub router, interface ge-0/0/1.
8. The hub PE router installs the routes learned from the hub CE router into its hub VRF table, which is associated with interface ge-0/0/1.
9. The hub PE router checks the VRF export policy associated with interface ge-0/0/1.0 and announces all routes that match to all spokes after adding the hub target community.

Figure 64 on page 667 illustrates how routes are distributed from this spoke router to the other spoke CE router, Router CE2. The same path is followed if you issue a **tracert** command from Router CE1 to Router CE2.

The final section in this example, “[Hub-and-Spoke VPN Configuration Summarized by Router](#)” on page 675, consolidates the statements needed to configure VPN functionality for each of the service provider routers shown in [Figure 63 on page 666](#).

Figure 64: Route Distribution Between Two Spoke Routers



The following sections explain how to configure the VPN functionality for a hub-and-spoke topology on the hub-and-spoke PE routers. The CE routers do not have any information about the VPN, so you configure them normally.

- [Enabling an IGP on the Hub-and-Spoke PE Routers on page 667](#)
- [Configuring LDP on the Hub-and-Spoke PE Routers on page 668](#)
- [Configuring IBGP on the PE Routers on page 668](#)
- [Configuring VPN Routing Instances on the Hub-and-Spoke PE Routers on page 669](#)
- [Configuring VPN Policy on the PE Routers on page 672](#)
- [Hub-and-Spoke VPN Configuration Summarized by Router on page 675](#)

## Enabling an IGP on the Hub-and-Spoke PE Routers

To allow the hub-and-spoke PE routers to exchange routing information, you must configure an IGP on all these routers or you must configure static routes. You configure the IGP on the master instance of the routing protocol process (`rp`) (that is, at the **[edit protocols]** hierarchy level), not within the routing instance (that is, not at the **[edit routing-instances]** hierarchy level).

You configure the IGP in the standard way. This configuration example does not include this portion of the configuration.

In the route distribution in a hub-and-spoke topology, if the protocol used between the CE and PE routers at the hub site is BGP, the hub CE router announces all routes received from the hub PE router and the spoke routers back to the hub PE router and all the spoke routers. This means that the hub-and-spoke PE routers receive routes that contain their AS number. Normally, when a route contains this information, it indicates that a routing loop has occurred and the router rejects the routes. However, for the VPN configuration to work, the hub PE router and the spoke routers must accept these routes. To enable this, include the **loops** option when configuring the AS at the **[edit routing-options]**

hierarchy level on the hub PE router and all the spoke routers. For this example configuration, you specify a value of 1. You can specify a number from 0 through 10.

```
[edit routing-options]
autonomous-system as-number loops 1;
```

## Configuring LDP on the Hub-and-Spoke PE Routers

Configure LDP on the interfaces between the hub-and-spoke PE routers that participate in the VPN.

On hub PE Router D, configure LDP:

```
[edit protocols]
ldp {
  interface so-1/0/0.0;
  interface t3-1/1/0.0;
}
```

On spoke PE Router E, configure LDP:

```
[edit protocols]
ldp {
  interface fe-0/1/2.0;
}
```

On spoke PE router Router F, configure LDP:

```
[edit protocols]
ldp {
  interface fe-1/0/0.0;
}
```

## Configuring IBGP on the PE Routers

On the hub-and-spoke PE routers, configure an IBGP session with the following properties:

- VPN family—To indicate that the IBGP session is for the VPN, include the **family inet-vpn** statement.
- Loopback address—Include the **local-address** statement, specifying the local PE router's loopback address. The IBGP session for VPNs runs through the loopback address. You must also configure the **lo0** interface at the **[edit interfaces]** hierarchy level. The example does not include this part of the router's configuration.
- Neighbor address—Include the **neighbor** statement. On the hub router, specify the IP address of each spoke PE router, and on the spoke router, specify the address of the hub PE router.

For the hub router, you configure an IBGP session with each spoke, and for each spoke router, you configure an IBGP session with the hub. There are no IBGP sessions between the two spoke routers.

On hub Router D, configure IBGP. The first **neighbor** statement configures an IBGP session to spoke Router E, and the second configures a session to spoke Router F.



```
[edit protocols]
bgp {
  group Hub-to-Spokes {
    type internal;
    local-address 10.255.14.174;
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.14.180;
    neighbor 10.255.14.182;
  }
}
```

On spoke Router E, configure an IBGP session to the hub router:

```
[edit protocols]
bgp {
  group Spoke-E-to-Hub {
    type internal;
    local-address 10.255.14.180;
    neighbor 10.255.14.174 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
```

On spoke Router F, configure an IBGP session to the hub router:

```
[edit protocols]
bgp {
  group Spoke-F-to-Hub {
    type internal;
    local-address 10.255.14.182;
    neighbor 10.255.14.174 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
```

## Configuring VPN Routing Instances on the Hub-and-Spoke PE Routers

For the hub PE router to be able to distinguish between packets going to and coming from the spoke PE routers, you must configure it with two routing instances:

- One routing instance (in this example, **Spokes-to-Hub-CE**) is associated with the interface that carries packets from the hub PE router to the hub CE router (in this example, interface **ge-0/0/0.0**). Its VRF table contains the routes being announced by the spoke PE routers and the hub PE router to the hub CE router.
- The second routing instance (in this example, **Hub-CE-to-Spokes**) is associated with the interface that carries packets from the hub CE router to the hub PE router (in this

example, interface **ge-0/0/1.0**). Its VRF table contains the routes being announced from the hub CE router to the hub-and-spoke PE routers.

On each spoke router, you must configure one routing instance.

You must define the following in the routing instance:

- Route distinguisher, which is used to distinguish the addresses in one VPN from those in another VPN.
- Instance type of **vrf**, which creates the VRF table on the PE router.
- Interfaces that are part of the VPN and that connect the PE routers to their CE routers.
- VRF import and export policies. Both import policies must include reference to a community. Otherwise, when you try to commit the configuration, the commit fails. (The exception to this is if the import policy contains only a **then reject** statement.) In the VRF export policy, spoke PE routers attach the spoke target community.
- Routing between the PE and CE routers, which is required for the PE router to distribute VPN-related routes to and from connected CE routers. You can configure a routing protocol—BGP, OSPF, or RIP—or you can configure static routing.

For a hub-and-spoke topology, you must configure different policies in each routing instance on the hub CE router. For the routing instance associated with the interface that carries packets from the hub PE router to the hub CE router (in this example, **Spokes-to-Hub-CE**), the import policy must accept all routes received on the IBGP session between the hub-and-spoke PE routers, and the export policy must reject all routes received from the hub CE router. For the routing instance associated with the interface that carries packets from the hub CE router to the hub PE router (in this example, **Hub-CE-to-Spokes**), the import policy must reject all routes received from the spoke PE routers, and the export policy must export to all the spoke routers.

On hub PE Router D, configure the following routing instances. Router D uses OSPF to distribute routes to and from the hub CE router.

```
[edit]
routing-instance {
  Spokes-to-Hub-CE {
    instance-type vrf;
    interface ge-0/0/0.0;
    route-distinguisher 10.255.1.174:65535;
    vrf-import spoke;
    vrf-export null;
    protocols {
      ospf {
        domain-id disable;
        export redistribute-vpn;
        domain-vpn-tag 0;
        area 0.0.0.0 {
          interface ge-0/0/0;
        }
      }
    }
  }
}
```

```

Hub-CE-to-Spokes {
  instance-type vrf;
  interface ge-0/0/1.0;
  route-distinguisher 10.255.1.174:65534;
  vrf-import null;
  vrf-export hub;
  protocols {
    ospf {
      export redistribute-vpn;
      area 0.0.0.0 {
        interface ge-0/0/1.0;
      }
    }
  }
}

```

On spoke PE Router E, configure the following routing instances. Router E uses OSPF to distribute routes to and from spoke CE Router CE1.

```

[edit]
routing-instance {
  Spoke-E-to-Hub {
    instance-type vrf;
    interface fe-0/1/0.0;
    route-distinguisher 10.255.14.80:65035;
    vrf-import hub;
    vrf-export spoke;
    protocols {
      ospf {
        export redistribute-vpn;
        area 0.0.0.0 {
          interface fe-0/1/0.0;
        }
      }
    }
  }
}

```

On spoke PE Router F, configure the following routing instances. Router F uses OSPF to distribute routes to and from spoke CE Router CE2.

```

[edit]
routing-instance {
  Spoke-F-to-Hub {
    instance-type vrf;
    interface fe-1/0/1.0;
    route-distinguisher 10.255.14.182:65135;
    vrf-import hub;
    vrf-export spoke;
    protocols {
      ospf {
        export redistribute-vpn;
        area 0.0.0.0 {
          interface fe-1/0/1.0;
        }
      }
    }
  }
}

```

```

    }
  }
}

```

## Configuring VPN Policy on the PE Routers

You must configure VPN import and export policies on each of the hub-and-spoke PE routers so that they install the appropriate routes in the VRF tables, which they use to forward packets within each VPN.

On the spoke routers, you define policies to exchange routes with the hub router.

On the hub router, you define policies to accept routes from the spoke PE routers and distribute them to the hub CE router, and vice versa. The hub PE router has two VRF tables:

- **Spoke-to-hub VRF table**—Handles routes received from spoke routers and announces these routes to the hub CE router. For this VRF table, the import policy must check that the spoke target name is present and that the route was received from the IBGP session between the hub PE and the spoke PE routers. This VRF table must not export any routes, so its export policy should reject everything.
- **Hub-to-spoke VRF table**—Handles routes received from the hub CE router and announces them to the spoke routers. For this VRF table, the export policy must add the hub target community. This VRF table must not import any routes, so its import policy should reject everything.

In the VPN policy, you also configure the VPN target communities.

On hub PE Router D, configure the following policies to apply to the VRF tables:

- **spoke**—Accepts routes received from the IBGP session between it and the spoke PE routers that contain the community target **spoke**, and rejects all other routes.
- **hub**—Adds the community target hub to all routes received from OSPF (that is, from the session between it and the hub CE router). It rejects all other routes.
- **null**—Rejects all routes.
- **redistribute-vpn**—Redistributes OSPF routes to neighbors within the routing instance.

```

[edit]
policy-options {
  policy-statement spoke {
    term a {
      from {
        protocol bgp;
        community spoke;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
}

```

```

policy-statement hub {
  term a {
    from protocol ospf;
    then {
      community add hub;
      accept;
    }
  }
  term b {
    then reject;
  }
}
policy-statement null {
  then reject;
}
policy-statement redistribute-vpn {
  term a {
    from protocol bgp;
    then accept;
  }
  term b {
    then reject;
  }
}
community hub members target:65535:1;
community spoke members target:65535:2;
}

```

To apply the VRF policies on Router D, include the **vrf-export** and **vrf-import** statements when you configure the routing instances:

```

[edit]
routing-instance {
  Spokes-to-Hub-CE {
    vrf-import spoke;
    vrf-export null;
  }
  Hub-CE-to-Spokes {
    vrf-import null;
    vrf-export hub;
  }
}
}

```

On spoke PE Router E and Router F, configure the following policies to apply to the VRF tables:

- **hub**—Accepts routes received from the IBGP session between it and the hub PE routers that contain the community target **hub**, and rejects all other routes.
- **spoke**—Adds the community target spoke to all routes received from OSPF (that is, from the session between it and the hub CE router) rejects all other routes.
- **redistribute-vpn**—Redistributes OSPF routes to neighbors within the routing instance.

On spoke PE Router E and Router F, configure the following VPN import and export policies:

```
[edit]
policy-options {
  policy-statement hub {
    term a {
      from {
        protocol bgp;
        community hub;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement spoke {
    term a {
      from protocol ospf;
      then {
        community add spoke;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  policy-statement redistribute-vpn {
    term a {
      from protocol bgp;
      then accept;
    }
    term b {
      then reject;
    }
  }
  community hub members target:65535:1;
  community spoke members target 65535:2;
}
```

To apply the VRF policies on the spoke routers, include the **vrf-export** and **vrf-import** statements when you configure the routing instances:

```
[edit]
routing-instance {
  Spoke-E-to-Hub {
    vrf-import hub;
    vrf-export spoke;
  }
}
[edit]
routing-instance {
  Spoke-F-to-Hub {
    vrf-import hub;
    vrf-export spoke;
  }
}
```

## Hub-and-Spoke VPN Configuration Summarized by Router

### Router D (Hub PE Router)

<b>Routing Instance for Distributing Spoke Routes to Hub CE</b>	<pre>Spokes-to-Hub-CE {   instance-type vrf;   interface fe-0/0/0.0;   route-distinguisher 10.255.1.174:65535;   vrf-import spoke;   vrf-export null; }</pre>
<b>Instance Routing Protocol</b>	<pre>protocols {   ospf {     domain-id disable;     domain-vpn-tag 0;     export redistribute-vpn;     area 0.0.0.0 {       interface ge-0/0/0.0;     }   } }</pre>
<b>Routing Instance for Distributing Hub CE Routes to Spokes</b>	<pre>Hub-CE-to-Spokes {   instance-type vrf;   interface ge-0/0/1.0;   route-distinguisher 10.255.1.174:65534;   vrf-import null;   vrf-export hub; }</pre>
<b>Routing Instance Routing Protocols</b>	<pre>protocols {   ospf {     export redistribute-vpn;     area 0.0.0.0 {       interface ge-0/0/1.0;     }   } }</pre>
<b>Routing Options (Master Instance)</b>	<pre>routing-options {   autonomous-system 1 loops 1; }</pre>
<b>Protocols (Master Instance)</b>	<pre>protocols { }</pre>
<b>Enable LDP</b>	<pre>ldp {   interface so-1/0/0.0;   interface t3-1/1/0.0; }</pre>

**Configure IBGP**

```
bgp {  
  group Hub-to-Spokes {  
    type internal;  
    local-address 10.255.14.174;  
    family inet-vpn {  
      unicast;  
    }  
    neighbor 10.255.14.180;  
    neighbor 10.255.14.182;  
  }  
}
```

**Configure VPN Policy**

```
policy-options {  
  policy-statement spoke {  
    term a {  
      from {  
        protocol bgp;  
        community spoke;  
      }  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement hub {  
    term a {  
      from protocol ospf;  
      then {  
        community add hub;  
        accept;  
      }  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement null {  
    then reject;  
  }  
  policy-statement redistribute-vpn {  
    term a {  
      from protocol bgp;  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  community hub members target:65535:1;  
  community spoke members target:65535:2;  
}
```



### Router E (Spoke PE Router)

<b>Routing Instance</b>	<pre> routing-instance {   Spoke-E-to-Hub {     instance-type vrf;     interface fe-0/1/0.0;     route-distinguisher 10.255.14.80:65035;     vrf-import hub;     vrf-export spoke;   } } </pre>
<b>Instance Routing Protocol</b>	<pre> protocols {   ospf {     export redistribute-vpn;     area 0.0.0.0 {       interface fe-0/1/0.0;     }   } } </pre>
<b>Routing Options (Master Instance)</b>	<pre> routing-options {   autonomous-system 1 loops 1; } </pre>
<b>Protocols (Master Instance)</b>	<pre> protocols { } </pre>
<b>Enable LDP</b>	<pre> ldp {   interface fe-0/1/2.0; } </pre>
<b>Configure IBGP</b>	<pre> bgp {   group Spoke-E-to-Hub {     type internal;     local-address 10.255.14.180;     neighbor 10.255.14.174 {       family inet-vpn {         unicast;       }     }   } } </pre>
<b>Configure VPN Policy</b>	<pre> policy-options {   policy-statement hub {     term a {       from {         protocol bgp;         community hub;       }     }   } } </pre>

```

        then accept;
    }
    term b {
        then reject;
    }
}
policy-statement spoke {
    term a {
        from protocol ospf;
        then {
            community add spoke;
            accept;
        }
    }
    term b {
        then reject;
    }
}
policy-statement redistribute-vpn {
    term a {
        from protocol bgp;
        then accept;
    }
    term b {
        then reject;
    }
}
community hub members target:65535:1;
community spoke members target:65535:2;
}

```

### Router F (Spoke PE Router)

<b>Routing Instance</b>	<pre> routing-instance {     Spoke-F-to-Hub {         instance-type vrf;         interface fe-1/0/1.0;         route-distinguisher 10.255.14.182:65135;         vrf-import hub;         vrf-export spoke;     } } </pre>
<b>Instance Routing Protocol</b>	<pre> protocols {     ospf {         export redistribute-vpn;         area 0.0.0.0 {             interface fe-1/0/1.0;         }     } } </pre>

Routing Options (Master Instance)	<pre>routing-options {   autonomous-system 1 loops 1; }</pre>
Protocols (Master Instance)	<pre>protocols { }</pre>
Enable LDP	<pre>ldp {   interface fe-1/0/0.0; }</pre>
Configure IBGP	<pre>bgp {   group Spoke-F-to-Hub {     type internal;     local-address 10.255.14.182;     neighbor 10.255.14.174 {       family inet-vpn {         unicast;       }     }   } }</pre>
Configure VPN Policy	<pre>policy-options {   policy-statement hub {     term a {       from {         protocol bgp;         community hub;       }       then accept;     }     term b {       then reject;     }   }   policy-statement spoke {     term a {       from protocol ospf;       then {         community add spoke;         accept;       }     }     term b {       then reject;     }   }   policy-statement redistribute-vpn {     term a {       from {         protocol bgp;       }     }   } }</pre>

```
        then accept;
    }
    term b {
        then reject;
    }
}
community hub members target:65535:1;
community spoke members target:65535:2;
}
```

# Configuring Overlapping VPNs

- [Configuring Overlapping VPNs Using Routing Table Groups on page 681](#)
- [Configuring Overlapping VPNs Using Automatic Route Export on page 692](#)

## Configuring Overlapping VPNs Using Routing Table Groups

---

In Layer 3 VPNs, a CE router is often a member of more than one VPN. This example illustrates how to configure PE routers that support CE routers that support multiple VPNs. Support for this type of configuration uses a Junos OS feature called routing table groups (sometimes also called routing information base [RIB] groups), which allows a route to be installed into several routing tables. A routing table group is a list of routing tables into which the protocol should install its routes.

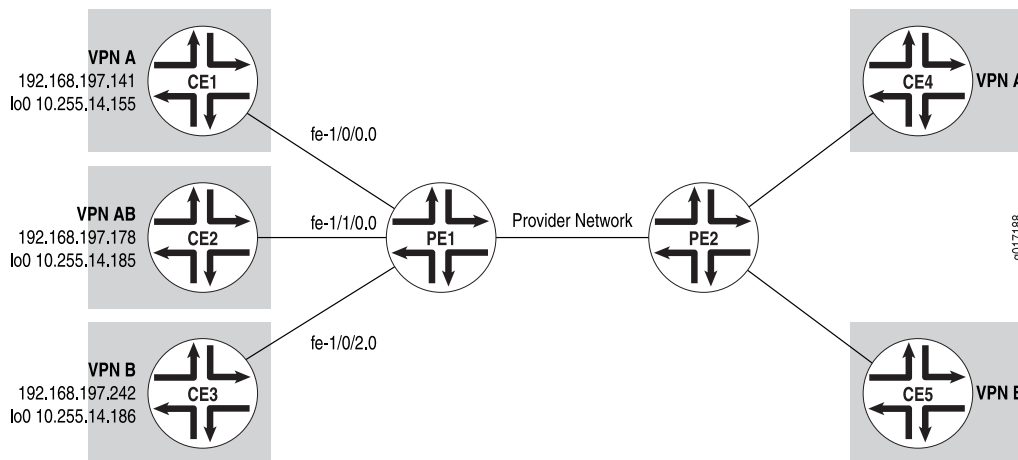
You define routing table groups at the **[edit routing-options]** hierarchy level for the default instance. You cannot configure routing table groups at the **[edit routing-instances routing-options]** hierarchy level; doing so results in a commit error.

After you define a routing table group, it can be used by multiple protocols. You can also apply routing table groups to static routing. The configuration examples in this section include both types of configurations.

[Figure 65 on page 682](#) illustrates the topology for the configuration example in this section. The configurations in this section illustrate local connectivity between CE routers connected to the same PE router. If Router PE1 were connected only to Router CE2 (VPN AB), there would be no need for any extra configuration. The configuration statements in the sections that follow enable VPN AB Router CE2 to communicate with VPN A Router CE1 and VPN B Router CE3, which are directly connected to Router PE1. VPN routes that originate from the remote PE routers (the PE2 router in this case) are placed in a global Layer 3 VPN routing table (bgp.l3vpn.inet.0), and routes with appropriate route targets are imported into the routing tables as dictated by the VRF import policy configuration. The goal is to be able to choose routes from individual VPN routing tables that are locally populated.

Router PE1 is where all the filtering and configuration modification takes place. Therefore only VPN configurations for PE1 are shown. The CE routers do not have any information about the VPN, so you can configure them normally.

Figure 65: Example of an Overlapping VPN Topology



The following sections explain several ways to configure overlapping VPNs.

The following sections illustrate different scenarios for configuring overlapping VPNs, depending on the routing protocol used between the PE and CE routers. For all of these examples, you need to configure routing table groups.

- [Configuring Routing Table Groups on page 682](#)
- [Configuring Static Routes Between the PE and CE Routers on page 683](#)
- [Configuring BGP Between the PE and CE Routers on page 688](#)
- [Configuring OSPF Between the PE and CE Routers on page 689](#)
- [Configuring Static, BGP, and OSPF Routes Between PE and CE Routers on page 691](#)

## Configuring Routing Table Groups

In this example, routing table groups are common in the four configuration scenarios. The routing table groups are used to install routes (including interface, static, OSPF, and BGP routes) into several routing tables for the default and other instances. In the routing table group definition, the first routing table is called the primary routing table. (Normally, the primary routing table is the table into which the route would be installed if you did not configure routing table groups. The other routing tables are called secondary routing tables.)

The routing table groups in this configuration install routes as follows:

- **vpn-a-vpnab** installs routes into routing tables VPN-A.inet.0 and VPN-AB.inet.0.
- **vpn-b-vpnab** installs routes into routing tables VPN-B.inet.0 and VPN-AB.inet.0.
- **vpnab-vpna\_and\_vpnb** installs routes into routing tables VPN-AB.inet.0, VPN-A.inet.0, and VPN-B.inet.0.

Configure the routing table groups:

```
[edit]
routing-options {
```

```

rib-groups {
  vpna-vpnab {
    import-rib [ VPN-A.inet.0 VPN-AB.inet.0 ];
  }
  vpnb-vpnab {
    import-rib [ VPN-B.inet.0 VPN-AB.inet.0 ];
  }
  vpnab-vpna_and_vpnb {
    import-rib [ VPN-AB.inet.0 VPN-A.inet.0 VPN-B.inet.0 ];
  }
}
}

```

## Configuring Static Routes Between the PE and CE Routers

To configure static routing between the PE1 router and the CE1, CE2, and CE3 routers, you must configure routing instances for VPN A, VPN B, and VPN AB (you configure static routing under each instance):

- [Configuring the Routing Instance for VPN A on page 683](#)
- [Configuring the Routing Instance for VPN AB on page 684](#)
- [Configuring the Routing Instance for VPN B on page 684](#)
- [Configuring VPN Policy on page 685](#)

### Configuring the Routing Instance for VPN A

On Router PE1, configure VPN A:

```

[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.14.175:3;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      interface-routes {
        rib-group inet vpna-vpnab;
      }
      static {
        route 10.255.14.155/32 next-hop 192.168.197.141;
        route 10.255.14.185/32 next-hop 192.168.197.178;
      }
    }
  }
}
}

```

The **interface-routes** statement installs VPN A's interface routes into the routing tables defined in the routing table group **vpna-vpnab**.

The **static** statement configures the static routes that are installed in the VPN-A.inet.0 routing table. The first static route is for Router CE1 (VPN A) and the second is for Router CE2 (in VPN AB).

Next hop **192.168.197.178** is not in VPN A. Route **10.255.14.185/32** cannot be installed in VPN-A.inet.0 unless interface routes from routing instance VPN AB are installed in this routing table. Including the **interface-routes** statements in the VPN AB configuration provides this next hop. Similarly, including the **interface-routes** statement in the VPN AB configuration installs **192.168.197.141** into VPN-AB.inet.0.

---

### Configuring the Routing Instance for VPN AB

On Router PE1, configure VPN AB:

```
[edit]
routing instances {
  VPN-AB {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpnab-import;
    vrf-export vpnab-export;
    routing-options {
      interface-routes {
        rib-group vpnab-vpna_and_vpnb;
      }
      static {
        route 10.255.14.185/32 next-hop 192.168.197.178;
        route 10.255.14.155/32 next-hop 192.168.197.141;
        route 10.255.14.186/32 next-hop 192.168.197.242;
      }
    }
  }
}
```

In this configuration, the following static routes are installed in the VPN-AB.inet.0 routing table:

- **10.255.14.185/32** is for Router CE2 (in VPN AB)
- **10.255.14.155/32** is for Router CE1 (in VPN A)
- **10.255.14.186/32** is for Router CE3 (in VPN B)

Next hops **192.168.197.141** and **192.168.197.242** do not belong to VPN AB. Routes **10.255.14.155/32** and **10.255.14.186/32** cannot be installed in VPN-AB.inet.0 unless interface routes from VPN A and VPN B are installed in this routing table. The interface route configurations in VPN A and VPN B routing instances provide these next hops.

---

### Configuring the Routing Instance for VPN B

On Router PE1, configure VPN B:

```
[edit]
routing instances {
  VPN-B {
    instance-type vrf;
    interface fe-1/0/2.0;
    route-distinguisher 10.255.14.175:10;
    vrf-import vpnb-import;
```



```

vrf-export vpnb-export;
routing-options {
  interface-routes {
    rib-group inet vpnb-vpnab;
  }
  static {
    route 10.255.14.186/32 next-hop 192.168.197.242;
    route 10.255.14.185/32 next-hop 192.168.197.178;
  }
}
}
}

```

When you configure the routing instance for VPN B, these static routes are placed in VPNB.inet.0:

- **10.255.14.186/32** is for Router CE3 (in VPN B)
- **10.255.14.185/32** is for Router CE2 (in VPN AB)

Next hop **192.168.197.178** does not belong to VPN B. Route **10.255.14.185/32** cannot be installed in VPN-B.inet.0 unless interface routes from VPN AB are installed in this routing table. The interface route configuration in VPN AB provides this next hop.

### Configuring VPN Policy

The **vrf-import** and **vrf-export** policy statements that you configure for overlapping VPNs are the same as policy statements for regular VPNs, except that you include the **from interface** statement in each VRF export policy. This statement forces each VPN to announce only those routes that originated from that VPN. For example, VPN A has routes that originated in VPN A and VPN AB. If you do not include the **from interface** statement, VPN A announces its own routes as well as VPN AB's routes, so the remote PE router receives multiple announcements for the same routes. Including the **from interface** statement restricts each VPN to announcing only the routes it originated and allows you to filter out the routes imported from other routing tables for local connectivity.

In this configuration example, the **vpnab-import** policy accepts routes from VPN A, VPN B, and VPN AB. The **vpna-export** policy exports only routes that originate in VPN A. Similarly, the **vpnb-export** and **vpnab-export** policies export only routes that originate within the respective VPNs.

On Router PE1, configure the following VPN import and export policies:

```

[edit]
policy-options {
  policy-statement vpna-import {
    term a {
      from {
        protocol bgp;
        community VPNA-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
}

```

```
    }  
  }  
  policy-statement vpnb-import {  
    term a {  
      from {  
        protocol bgp;  
        community VPNB-comm;  
      }  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement vpnab-import {  
    term a {  
      from {  
        protocol bgp;  
        community [ VPNA-comm VPNB-comm ];  
      }  
      then accept;  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement vpna-export {  
    term a {  
      from {  
        protocol static;  
        interface fe-1/0/0.0;  
      }  
      then {  
        community add VPNA-comm;  
        accept;  
      }  
    }  
    term b {  
      then reject;  
    }  
  }  
  policy-statement vpnb-export {  
    term a {  
      from {  
        protocol static;  
        interface fe-1/0/2.0;  
      }  
      then {  
        community add VPNB-comm;  
        accept;  
      }  
    }  
    term b {  
      then reject;  
    }  
  }  
}
```

```

policy-statement vpnab-export {
  term a {
    from {
      protocol static;
      interface fe-1/1/0.0;
    }
    then {
      community add VPNB-comm;
      community add VPNA-comm;
      accept;
    }
  }
  term b {
    then reject;
  }
}
community VPNA-comm members target:69:1;
community VPNB-comm members target:69:2;
}

```

On Router PE1, apply the VPN import and export policies:

```

[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.14.175:3;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        rib-group vpna-vpnab;
        route 10.255.14.155/32 next-hop 192.168.197.141;
        route 10.255.14.185/32 next-hop 192.168.197.178;
      }
    }
  }
  VPN-AB {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpnab-import;
    vrf-export vpnab-export;
    routing-options {
      static {
        rib-group vpnab-vpna_and_vpnab;
        route 10.255.14.185/32 next-hop 192.168.197.178;
      }
    }
  }
  VPN-B {
    instance-type vrf;
    interface fe-1/0/2.0;
    route-distinguisher 10.255.14.175:10;
    vrf-import vpnb-import;
  }
}

```

```

vrf-export vpnb-export;
routing-options {
  static {
    rib-group vpnb-vpnab;
    route 10.255.14.186/32 next-hop 192.168.197.242;
  }
}
}
}

```

For VPN A, include the **routing-options** statement at the **[edit routing-instances routing-instance-name]** hierarchy level to install the static routes directly into the routing tables defined in the routing table group **vpnab-vpnab**. For VPN AB, the configuration installs the static route directly into the routing tables defined in the routing table group **vpnab-vpna** and **vpnab-vpnb**. For VPN B the configuration installs the static route directly into the routing tables defined in the routing table group **vpnb-vpnab**.

## Configuring BGP Between the PE and CE Routers

In this configuration example, the **vpna-site1** BGP group for VPN A installs the routes learned from the BGP session into the routing tables defined in the **vpnab-vpnab** routing table group. For VPN AB, the **vpnab-site1** group installs the routes learned from the BGP session into the routing tables defined in the **vpnab-vpna\_and\_vpnb** routing table group. For VPN B, the **vpnb-site1** group installs the routes learned from the BGP session into the routing tables defined in the **vpnb-vpnab** routing table group. Interface routes are not needed for this configuration.

The VRF import and export policies are similar to those defined in “[Configuring Static Routes Between the PE and CE Routers](#)” on page 683, except the export protocol is BGP instead of a static route. On all **vrf-export** policies, you use the **from protocol bgp** statement.

On Router PE1, configure BGP between the PE and CE routers:

```

[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.14.175:3;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group vpna-site1 {
          family inet {
            unicast {
              rib-group vpna-vpnab;
            }
          }
        }
        peer-as 1;
        neighbor 192.168.197.141;
      }
    }
  }
}

```

```

    }
  }
  VPN-AB {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpnab-import;
    vrf-export vpnab-export;
    protocols {
      bgp {
        group vpnab-site1 {
          family inet {
            unicast {
              rib-group vpnab-vpna_and_vpnb;
            }
          }
        }
        peer-as 9;
        neighbor 192.168.197.178;
      }
    }
  }
  VPN-B {
    instance-type vrf;
    interface fe-1/0/2.0;
    route-distinguisher 10.255.14.175:10;
    vrf-import vpnb-import;
    vrf-export vpnb-export;
    protocols {
      bgp {
        group vpnb-site1 {
          family inet {
            unicast {
              rib-group vpnb-vpnab;
            }
          }
        }
        neighbor 192.168.197.242 {
          peer-as 10;
        }
      }
    }
  }
}

```

## Configuring OSPF Between the PE and CE Routers

In this configuration example, routes learned from the OSPF session for VPN A are installed into the routing tables defined in the **vpna-vpnab** routing table group. For VPN AB, routes learned from the OSPF session are installed into the routing tables defined in the **vpnab-vpna\_and\_vpnb** routing table group. For VPN B, routes learned from the OSPF session are installed into the routing tables defined in the **vpnb-vpnab** routing table group.

The VRF import and export policies are similar to those defined in “[Configuring Static Routes Between the PE and CE Routers](#)” on page 683 and “[Configuring BGP Between the](#)

[PE and CE Routers](#) on page 688, except the export protocol is OSPF instead of BGP or a static route. Therefore, on all **vrf-export** policies, you use the **from protocol ospf** statement instead of the **from protocol <static | bgp>** statement.

On Router PE1, configure OSPF between the PE and CE routers:

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.14.175:3;
    vrf-import vpn-a-import;
    vrf-export vpn-a-export;
    protocols {
      ospf {
        rib-group vpn-a-vpnab;
        export vpn-a-import;
        area 0.0.0.0 {
          interface fe-1/0/0.0;
        }
      }
    }
  }
  VPN-AB {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpnab-import;
    vrf-export vpnab-export;
    protocols {
      ospf {
        rib-group vpnab-vpn-a_and_vpn-b;
        export vpnab-import;
        area 0.0.0.0 {
          interface fe-1/1/0.0;
        }
      }
    }
  }
  VPN-B {
    instance-type vrf;
    interface fe-1/0/2.0;
    route-distinguisher 10.255.14.175:10;
    vrf-import vpnb-import;
    vrf-export vpnb-export;
    protocols {
      ospf {
        rib-group vpnb-vpnab;
        export vpnb-import;
        area 0.0.0.0 {
          interface fe-1/0/2.0;
        }
      }
    }
  }
}
```

```
}
```

## Configuring Static, BGP, and OSPF Routes Between PE and CE Routers

This section shows how to configure the routes between the PE and CE routers by using a combination of static routes, BGP, and OSPF:

- The connection between Router PE1 and Router CE1 uses static routing.
- The connection between Router PE1 and Router CE2 uses BGP.
- The connection between Router PE1 and Router CE3 uses OSPF.

Here, the configuration for VPN AB also includes a static route to CE1.

On Router PE1, configure a combination of static routing, BGP, and OSPF between the PE and CE routers:

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.14.175:3;
    vrf-import vpna-import;
    vrf-export vpna-export;
    routing-options {
      static {
        rib-group vpna-vpnab;
        route 10.255.14.155/32 next-hop 192.168.197.141;
      }
    }
  }
  VPN-AB {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpnab-import;
    vrf-export vpnab-export;
    protocols {
      bgp {
        group vpnab-site1 {
          family inet {
            unicast {
              rib-group vpnab-vpna_and_vpnab;
            }
          }
          export to-vpnab-site1;
          peer-as 9;
          neighbor 192.168.197.178;
        }
      }
    }
  }
  VPN-B {
    instance-type vrf;
```

```
interface fe-1/0/2.0;
route-distinguisher 10.255.14.175:10;
vrf-import vpnb-import;
vrf-export vpnb-export;
protocols {
  ospf {
    rib-group vpnb-vpnab;
    export vpnb-import;
    area 0.0.0.1 {
      interface t3-0/3/3.0;
    }
  }
}
}
policy-options {
  policy-statement to-vpnab-site1 {
    term a {
      from protocol static;
      then accept;
    }
    term b {
      from protocol bgp;
      then accept;
    }
    term c {
      then reject;
    }
  }
}
```

---

## Configuring Overlapping VPNs Using Automatic Route Export

A problem with multiple routing instances is how to export routes between routing instances. You can accomplish this in Junos OS by configuring routing table groups for each routing instance that needs to export routes to other routing tables. For information about how to configure overlapping VPNs by using routing table groups, see [“Configuring Overlapping VPNs Using Routing Table Groups”](#) on page 681.

However, using routing table groups has limitations:

- Routing table group configuration is complex. You must define a unique routing table group for each routing instance that will export routes.
- You must also configure a unique routing table group for each protocol that will export routes.

To limit and sometimes eliminate the need to configure routing table groups in multiple routing instance topologies, you can use the functionality provided by the **auto-export** statement.

The **auto-export** statement is particularly useful for configuring overlapping VPNs—VPN configurations where more than one VRF routing instance lists the same community route target in its **vrf-import** policy. The **auto-export** statement finds out which routing



tables to export routes from and import routes to by examining the existing policy configuration.

The **auto-export** statement automatically exports routes between the routing instances referencing a given route target community. When the **auto-export** statement is configured, a VRF target tree is constructed based on the **vrf-import** and **vrf-export** policies configured on the system. If a routing instance references a route target in its **vrf-import** policy, the route target is added to the import list for the target. If it references a specific route target in its **vrf-export** policy, the route target is added to the export list for that target. Route targets where there is a single importer that matches a single exporter or with no importers or exporters are ignored.

Changes to routing tables that export route targets are tracked. When a route change occurs, the routing instance's **vpn-export** policy is applied to the route. If it is allowed, the route is imported to all the import tables (subject to the **vrf-import** policy) of the route targets set by the export policy.

The sections that follow describe how to configure overlapping VPNs by using the **auto-export** statement for inter-instance export in addition to routing table groups:

- [Configuring Overlapping VPNs with BGP and Automatic Route Export on page 693](#)
- [Configuring Overlapping VPNs and Additional Tables on page 694](#)
- [Configuring Automatic Route Export for All VRF Instances on page 695](#)

## Configuring Overlapping VPNs with BGP and Automatic Route Export

The following example provides the configuration for an overlapping VPN where BGP is used between the PE and CE routers.

Configure routing instance **VPN-A**:

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.14.175:3;
    vrf-import vpn-a-import;
    vrf-export vpn-a-export;
    routing-options {
      auto-export;
    }
    protocols {
      bgp {
        group vpn-a-site1 {
          peer-as 1;
          neighbor 192.168.197.141;
        }
      }
    }
  }
}
```

Configure routing instance **VPN-AB**:

```
[edit]
routing-instances {
  VPN-AB {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpnab-import;
    vrf-export vpnab-export;
    routing-options {
      auto-export;
    }
    protocols {
      bgp {
        group vpnab-site1 {
          peer-as 9;
          neighbor 192.168.197.178;
        }
      }
    }
  }
}
```

For this configuration, the **auto-export** statement replaces the functionality that was provided by a routing table group configuration. However, sometimes additional configuration is required.

Since the **vrf-import** policy and the **vrf-export** policy from which the **auto-export** statement deduces the import and export matrix are configured on a per-instance basis, you must be able to enable or disable them for unicast and multicast, in case multicast network layer reachability information (NLRI) is configured.

## Configuring Overlapping VPNs and Additional Tables

You might need to use the **auto-export** statement between overlapping VPNs but require that a subset of the routes learned from a VRF table be installed into the inet.0 table or in routing-instance.inet.2.

To support this type of scenario, where not all of the information needed is present in the **vrf-import** and **vrf-export** policies, you configure an additional list of routing tables by using an additional routing table group.

To add routes from **VPN-A** and **VPN-AB** to inet.0 in the example described, you need to include the following additional configuration statements:

Configure the routing options:

```
[edit]
routing-options {
  rib-groups {
    inet-access {
      import-rib inet.0;
    }
  }
}
```

```
}
```

Configure routing instance **VPN-A**:

```
[edit]
routing-instances {
  VPN-A {
    routing-options {
      auto-export {
        family inet {
          unicast {
            rib-group inet-access;
          }
        }
      }
    }
  }
}
```

Configure routing instance **VPN-AB**:

```
[edit]
routing-instances {
  VPN-AB {
    routing-options {
      auto-export {
        family inet {
          unicast {
            rib-group inet-access;
          }
        }
      }
    }
  }
}
```

Routing table groups are used in this configuration differently from how they are generally used in Junos OS. Routing table groups normally require that the exporting routing table be referenced as the primary import routing table in the routing table group. For this configuration, the restriction does not apply. The routing table group functions as an additional list of tables to which to export routes.

## Configuring Automatic Route Export for All VRF Instances

The following configuration allows you to configure the **auto-export** statement for all of the routing instances in a configuration group:

```
[edit]
groups {
  vrf-export-on {
    routing-instances {
      <*> {
        routing-options {
          auto-export;
        }
      }
    }
  }
}
```

```
    }  
  }  
  apply-groups vrf-export-on;
```

## PART 6

# Configuring IPSec VPNs in Layer 3 VPNs

- [Configuring IPSec Tunnels on page 699](#)
- [Configuring ES Tunnel Interfaces for Layer 3 VPNs on page 703](#)
- [Configuring GRE Tunnel Interfaces for Layer 3 VPNs on page 711](#)
- [Configuring Next-Hop-Based Dynamic Tunnels for Layer 3 VPNs on page 727](#)



# Configuring IPsec Tunnels

- [Configuring IPsec Tunnels Instead of MPLS LSPs Between PE Routers in Layer 3 VPNs on page 699](#)

## Configuring IPsec Tunnels Instead of MPLS LSPs Between PE Routers in Layer 3 VPNs

A conventional Layer 3 BGP/MPLS VPN requires the configuration of MPLS label-switched paths (LSPs) between the PE routers. When a PE router receives a packet from a CE router, it performs a lookup in a specific VRF table for the IP destination address and obtains a corresponding MPLS label stack. The label stack is used to forward the packet to the egress PE router, where the bottom label is removed and the packet is forwarded to the specified CE router.

You can provide Layer 3 BGP/MPLS VPN service without an MPLS backbone. Instead of configuring MPLS LSPs between the PE routers, you configure GRE and IPsec tunnels between the PE routers. The MPLS information for the VPN (the VPN label) is encapsulated within an IP header and an IPsec header. The source address of the IP header is the address of the ingress PE router. The destination address has the BGP next hop, the address of the egress PE router.



**NOTE:** The IPsec tunnel requires the use of an ES PIC. The GRE tunnel requires the use of a Tunnel Services PIC.

To configure IPsec between PE routers, follow these steps:

1. Configure an IPsec tunnel between the PE routers. The source address is that of the ingress PE router, and the destination address is that of the egress PE router:

```
es-interface-name {
  unit unit-number {
    tunnel {
      source source-address;
      destination destination-address;
    }
    family inet {
      ipsec-sa sa-esp-dynamic;
      address address;
    }
    family mpls;
```

```
    }  
  }
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

2. Configure IPsec on the PE router. For information about how to configure IPsec, see the *Junos OS Administration Library*.
3. Configure a GRE tunnel between the PE routers. Again, the source address is that of the ingress PE router, and the destination address is that of the egress PE router:

```
gr-interface-name {  
  unit unit-number {  
    family inet {  
      address address;  
    }  
    family mpls;  
    tunnel {  
      source source-address;  
      destination destination-address;  
    }  
  }  
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

4. Configure BGP between the PE routers:

```
bgp {  
  group pe {  
    type internal;  
    local-address local-address;  
    family inet {  
      unicast;  
    }  
    family inet-vpn {  
      unicast;  
    }  
    peer-as as-number;  
    neighbor address;  
  }  
}
```

You can include these statements at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

5. Configure the routing instance:



```

instance-type vrf;
interface interface-name;
route-distinguisher address;
vrf-import import-policy-name;
vrf-export export-policy-name;
protocols {
  bgp {
    group routing-instance-name {
      type external;
      peer-as as-number;
      as-override;
      neighbor address;
    }
  }
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

6. Configure the policy options:

```

policy-statement import-policy-name {
  term 1 {
    from {
      protocol bgp;
      community community-name;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
policy-statement export-policy-name {
  term 1 {
    from protocol [ bgp direct ];
    then {
      community add community-name;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}
community community-name members target:target;

```

You can include these statements at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

7. Configure routing table groups to enable VPN route resolution in the inet.3 routing table:

```
interface-routes {  
    rib-group inet if-rib;  
}  
rib inet.3 {  
    static {  
        route BGP-address-for-remote-PE next-hop gre-interface-name;  
    }  
}  
rib-groups {  
    if-rib {  
        import-rib [ inet.0 inet.3 ];  
    }  
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

# Configuring ES Tunnel Interfaces for Layer 3 VPNs

- [Configuring an ES Tunnel Interface for Layer 3 VPNs on page 703](#)
- [Configuring an ES Tunnel Interface Between a PE and CE Router on page 705](#)

## Configuring an ES Tunnel Interface for Layer 3 VPNs

---

An ES tunnel interface allows you to configure an IP Security (IPsec) tunnel between the PE and CE routers of a Layer 3 VPN. The IPsec tunnel can include one or more hops.

The following sections explain how to configure an ES tunnel interface between the PE and CE routers of a Layer 3 VPN:

- [Configuring the ES Tunnel Interface on the PE Router on page 703](#)
- [Configuring the ES Tunnel Interface on the CE Router on page 704](#)

## Configuring the ES Tunnel Interface on the PE Router

To configure the ES tunnel interface on the PE router, include the **unit** statement:

```
unit logical-unit-number {  
  tunnel {  
    source source-address;  
    destination destination-address;  
  }  
  family inet {  
    address address;  
    ipsec-sa security-association-name;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name*]**

By default, the tunnel destination address is assumed to be in the default Internet routing table, inet.0. For IPsec tunnels using manual security association (SA), if the tunnel destination address is not in the default inet.0 routing table, you need to specify which

routing table to search for the tunnel destination address by configuring the **routing-instance** statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

```
unit logical-unit-number {
  tunnel {
    source address;
    destination address;
    routing-instance {
      destination routing-instance-name;
    }
    family inet {
      address address;
      ipsec-sa security-association-name;
    }
    family mpls;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]



**NOTE:** For IPsec tunnels using dynamic SA, the tunnel destination address must be in the default Internet routing table, inet.0.

To complete the ES tunnel interface configuration, include the **interface** statement for the ES interface under the appropriate routing instance:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

## Configuring the ES Tunnel Interface on the CE Router

To configure the ES tunnel interface on the CE router, include the **unit** statement:

```
unit 0 {
  tunnel {
    source address;
    destination address;
  }
  family inet {
    address address;
    ipsec-sa security-association-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name]`
- `[edit logical-systems logical-system-name interfaces interface-name]`


For more information about how to configure tunnel interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

For more information about how to configure IPsec interfaces, see the *Junos OS Administration Library*.

## Configuring an ES Tunnel Interface Between a PE and CE Router

This example shows how to configure an ES tunnel interface between a PE router and a CE router in a Layer 3 VPN. The network topology used in this example is shown in [Figure 66 on page 705](#).

Figure 66: ES Tunnel Interface (IPsec Tunnel)



To configure this example, you perform the steps in the following sections:

- [Configuring IPsec on Router PE1 on page 705](#)
- [Configuring the Routing Instance Without the Encapsulating Interface on page 706](#)
- [Configuring the Routing Instance with the Encapsulating Interface on page 707](#)
- [Configuring the ES Tunnel Interface on Router CE1 on page 708](#)
- [Configuring IPsec on Router CE1 on page 708](#)

## Configuring IPsec on Router PE1

Configure IP Security (IPsec) on Router PE1:

```
[edit security]
ipsec {
  security-association sa-esp-manual {
    mode tunnel;
    manual {
      direction bidirectional {
        protocol esp;
        spi 16000;
        authentication {
          algorithm hmac-md5-96;
          key ascii-text
            "$9$ABULT1heK87dsWLDk.P3nrevM7V24ZHkPaZ/tpOcSvWLNwgZUH";
        }
        encryption {
```

```
        algorithm des-cbc;
        key ascii-text "$9$/H8Q90lrvL7VKMZjHqQzcyleLN";
    }
}
}
```

## Configuring the Routing Instance Without the Encapsulating Interface

You can configure the routing instance on Router PE1 with or without the encapsulating interface (**t3-0/1/3** in this example). The following sections explain how to configure the routing instance without it:

- [Configuring the Routing Instance on Router PE1 on page 706](#)
- [Configuring the ES Tunnel Interface on Router PE1 on page 706](#)
- [Configuring the Encapsulating Interface for the ES Tunnel on page 707](#)

### Configuring the Routing Instance on Router PE1

---

Configure the routing instance on Router PE1:

```
[edit routing-instances]
vpna {
  instance-type vrf;
  interface es-1/2/0.0;
  route-distinguisher 10.255.14.174:1;
  vrf-import vpna-import;
  vrf-export vpna-export;
  protocols {
    bgp {
      group vpna {
        type external;
        peer-as 100;
        as-override;
        neighbor 10.49.2.1;
      }
    }
  }
}
```

### Configuring the ES Tunnel Interface on Router PE1

---

Configure the ES tunnel interface on Router PE1:

```
[edit interfaces es-1/2/0]
unit 0 {
  tunnel {
    source 192.168.197.249;
    destination 192.168.197.250;
  }
  family inet {
    address 10.49.2.2/30;
    ipsec-sa sa-esp-manual;
  }
}
```

```
}
```

### Configuring the Encapsulating Interface for the ES Tunnel

For this example, interface **t3-0/1/3** is the encapsulating interface for the ES tunnel. Configure interface **t3-0/1/3**:

```
[edit interfaces t3-0/1/3]
unit 0 {
  family inet {
    address 192.168.197.249/30;
  }
}
```

### Configuring the Routing Instance with the Encapsulating Interface

If the tunnel-encapsulating interface, **t3-0/1/3**, is also configured under the routing instance, you need to specify the routing instance name under the interface definition. The system uses this routing instance to search for the tunnel destination address for the IPsec tunnel using manual security association.

The following sections explain how to configure the routing instance with the encapsulating interface:

- [Configuring the Routing Instance on Router PE1 on page 707](#)
- [Configuring the ES Tunnel Interface on Router PE1 on page 708](#)
- [Configuring the Encapsulating Interface on Router PE1 on page 708](#)

### Configuring the Routing Instance on Router PE1

Configure the routing instance on Router PE1 (including the tunnel encapsulating interface):

```
[edit routing-instances]
vpna {
  instance-type vrf;
  interface es-1/2/0.0;
  interface t3-0/1/3.0;
  route-distinguisher 10.255.14.174:1;
  vrf-import vpna-import;
  vrf-export vpna-export;
  protocols {
    bgp {
      group vpna {
        type external;
        peer-as 100;
        as-override;
        neighbor 10.49.2.1;
      }
    }
  }
}
```

### Configuring the ES Tunnel Interface on Router PE1

---

Configure the ES tunnel interface on Router PE1:

```
[edit interfaces es-1/2/0]
unit 0 {
  tunnel {
    source 192.168.197.249;
    destination 192.168.197.250;
    routing-instance {
      destination vpna;
    }
  }
  family inet {
    address 10.49.2.2/30;
    ipsec-sa sa-esp-manual;
  }
}
```

### Configuring the Encapsulating Interface on Router PE1

---

Configure the encapsulating interface on Router PE1:

```
[edit interfaces t3-0/1/3]
unit 0 {
  family inet {
    address 192.168.197.249/30;
  }
}
```

### Configuring the ES Tunnel Interface on Router CE1

Configure the ES tunnel interface on Router CE1:

```
[edit interfaces es-1/2/0]
unit 0 {
  tunnel {
    source 192.168.197.250;
    destination 192.168.197.249;
  }
  family inet {
    address 10.49.2.1/30;
    ipsec-sa sa-esp-manual;
  }
}
```

### Configuring IPsec on Router CE1

Configure IPsec on Router CE1:

```
[edit security]
ipsec {
  security-association sa-esp-manual {
    mode tunnel;
    manual {
      direction bidirectional {
```



```
protocol esp;
spi 16000;
authentication {
    algorithm hmac-md5-96;
    key ascii-text
        "$9$ABULt1heK87dsWLDk.P3nrevM7V24ZHkPaZ/tpOcSvWLNwgZUH";
}
encryption {
    algorithm des-cbc;
    key ascii-text "$9$/H8Q90IyrvL7VKMZjHqQzcyleLN";
}
}
}
}
```



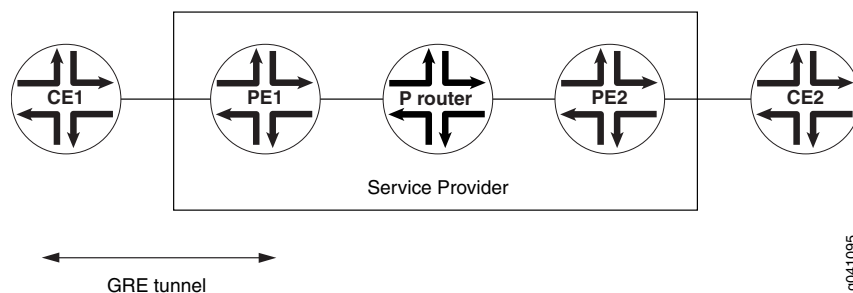
# Configuring GRE Tunnel Interfaces for Layer 3 VPNs

- [Configuring GRE Tunnels for Layer 3 VPNs on page 711](#)
- [Configuring a GRE Tunnel Interface Between PE Routers on page 714](#)
- [Configuring a GRE Tunnel Interface Between a PE and CE Router on page 721](#)

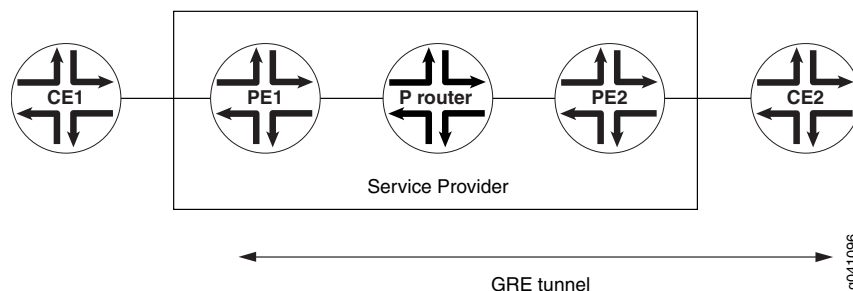
## Configuring GRE Tunnels for Layer 3 VPNs

Junos OS allows you to configure a generic routing encapsulation (GRE) tunnel between the PE and CE routers for a Layer 3 VPN. The GRE tunnel can have one or more hops. You can configure the tunnel from the PE router to a local CE router (as shown in [Figure 67 on page 711](#)) or to a remote CE router (as shown in [Figure 68 on page 711](#)).

*Figure 67: GRE Tunnel Configured Between the Local CE Router and the PE Router*



*Figure 68: GRE Tunnel Configured Between the Remote CE Router and the PE Router*



For more information about how to configure tunnel interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

You can configure the GRE tunnels manually or configure the Junos OS to instantiate GRE tunnels dynamically.

The following sections describe how to configure GRE tunnels manually and dynamically:

- [Configuring GRE Tunnels Manually Between PE and CE Routers on page 712](#)
- [Configuring GRE Tunnels Dynamically on page 713](#)

## Configuring GRE Tunnels Manually Between PE and CE Routers

You can manually configure a GRE tunnel between a PE router and either a local CE router or a remote CE router for a Layer 3 VPN as explained in the following sections:

- [Configuring the GRE Tunnel Interface on the PE Router on page 712](#)
- [Configuring the GRE Tunnel Interface on the CE Router on page 713](#)

---

### Configuring the GRE Tunnel Interface on the PE Router

You configure the GRE tunnel as a logical interface on the PE router. To configure the GRE tunnel interface, include the **unit** statement:

```
unit logical-unit-number {  
  tunnel {  
    source source-address;  
    destination destination-address;  
    routing-instance {  
      destination routing-instance-name;  
    }  
  }  
  family inet {  
    address address;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name*]**

As part of the GRE tunnel interface configuration, you need to include the following statements:

- **source *source-address***—Specify the source or origin of the GRE tunnel, typically the PE router.
- **destination *destination-address***—Specify the destination or end point of the GRE tunnel. The destination can be a Provider router, the local CE router, or the remote CE router.

By default, the tunnel destination address is assumed to be in the default Internet routing table, `inet.0`. If the tunnel destination address is not in `inet.0`, you need to specify which routing table to search for the tunnel destination address by configuring the **routing-instance** statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

- **destination *routing-instance-name***—Specify the name of the routing instance when configuring the GRE tunnel interface on the PE router.

To complete the GRE tunnel interface configuration, include the **interface** statement for the GRE interface under the appropriate routing instance:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

### Configuring the GRE Tunnel Interface on the CE Router

You can configure either the local or the remote CE router to act as the endpoint for the GRE tunnel.

To configure the GRE tunnel interface on the CE router, include the **unit** statement:

```
unit logical-unit-number {
  tunnel {
    source address;
    destination address;
  }
  family inet {
    address address;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name*]**

### Configuring GRE Tunnels Dynamically

When the router receives a VPN route to a BGP next hop address, but no MPLS path is available, a GRE tunnel can be dynamically generated to carry the VPN traffic across the BGP network. The GRE tunnel is generated and then its routing information is copied into the `inet.3` routing table. IPv4 routes are the only type of routes supported for dynamic GRE tunnels. Also, the routing platform must have a tunnel PIC.



**NOTE:** When configuring a dynamic GRE tunnel to a remote CE router, do not configure OSPF over the tunnel interface. It creates a routing loop forcing the router to take the GRE tunnel down. The router attempts to reestablish the GRE tunnel, but will be forced to take it down again when OSPF becomes active on the tunnel interface and discovers a route to the tunnel endpoint. This is not an issue when configuring static GRE tunnels to a remote CE router.

To generate GRE tunnels dynamically, include the **dynamic-tunnels** statement:

```
dynamic-tunnels tunnel-name {
    destination-networks prefix;
    source-address address;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

Specify the IPv4 prefix range (for example, 10/8 or 11.1/16) for the destination network by including the **destination-networks** statement. Only tunnels within the specified IPv4 prefix range are allowed to be initiated.

```
destination-networks prefix;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]

Specify the source address for the GRE tunnels by including the **source-address** statement. The source address specifies the address used as the source for the local tunnel endpoint. This could be any local address on the router (typically the router ID or the loopback address).

```
source-address address;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]

#### Related Documentation


- *Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks*

## Configuring a GRE Tunnel Interface Between PE Routers

This example shows how to configure a generic routing encapsulation (GRE) tunnel interface between PE routers to provide VPN connectivity. You can use this configuration to tunnel VPN traffic across a non-MPLS core network. The network topology used in

this example is shown in [Figure 69 on page 715](#). The P routers shown in this illustration do not run MPLS.

*Figure 69: PE Routers A and D Connected by a GRE Tunnel Interface*



For configuration information, see the following sections:

- [Configuring the Routing Instance on Router A on page 715](#)
- [Configuring the Routing Instance on Router D on page 716](#)
- [Configuring MPLS, BGP, and OSPF on Router A on page 716](#)
- [Configuring MPLS, BGP, and OSPF on Router D on page 716](#)
- [Configuring the Tunnel Interface on Router A on page 717](#)
- [Configuring the Tunnel Interface on Router D on page 717](#)
- [Configuring the Routing Options on Router A on page 718](#)
- [Configuring the Routing Options on Router D on page 718](#)
- [Configuration Summary for Router A on page 718](#)
- [Configuration Summary for Router D on page 720](#)

## Configuring the Routing Instance on Router A

Configure a routing instance on Router A:

```
[edit routing-instances]
gre-config {
  instance-type vrf;
  interface fe-1/0/0.0;
  route-distinguisher 10.255.14.176:69;
  vrf-import import-config;
  vrf-export export-config;
  protocols {
    ospf {
      export import-config;
      area 0.0.0.0 {
        interface all;
      }
    }
  }
}
```

## Configuring the Routing Instance on Router D

Configure a routing instance on Router D:

```
[edit routing-instances]
gre-config {
  instance-type vrf;
  interface fe-1/0/1.0;
  route-distinguisher 10.255.14.178:69;
  vrf-import import-config;
  vrf-export export-config;
  protocols {
    ospf {
      export import-config;
      area 0.0.0.0 {
        interface all;
      }
    }
  }
}
```

## Configuring MPLS, BGP, and OSPF on Router A

Although you do not need to configure MPLS on the P routers in this example, it is needed on the PE routers for the interface between the PE and CE routers and on the GRE interface (**gr-1/1/0.0**) linking the PE routers (Router A and Router D). Configure MPLS, BGP, and OSPF on Router A:

```
[edit protocols]
mpls {
  interface all;
}
bgp {
  group pe-to-pe {
    type internal;
    neighbor 10.255.14.178 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface all;
    interface gr-1/1/0.0 {
      disable;
    }
  }
}
```

## Configuring MPLS, BGP, and OSPF on Router D

Although you do not need to configure MPLS on the P routers in this example, it is needed on the PE routers for the interface between the PE and CE routers and on the GRE interface



(gr-1/1/0.0) linking the PE routers (Router D and Router A). Configure MPLS, BGP, and OSPF on Router D:

```
[edit protocols]
mpls {
  interface all;
}
bgp {
  group pe-to-pe {
    type internal;
    neighbor 10.255.14.176 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface gr-1/1/0.0 {
      disable;
    }
  }
}
```

## Configuring the Tunnel Interface on Router A

Configure the tunnel interface on Router A (the tunnel is unnumbered):

```
[edit interfaces interface-name]
unit 0 {
  tunnel {
    source 10.255.14.176;
    destination 10.255.14.178;
  }
  family inet;
  family mpls;
}
```

## Configuring the Tunnel Interface on Router D

Configure the tunnel interface on Router D (the tunnel is unnumbered):

```
[edit interfaces interface-name]
unit 0 {
  tunnel {
    source 10.255.14.178;
    destination 10.255.14.176;
  }
  family inet;
  family mpls;
}
```

```
}
```

## Configuring the Routing Options on Router A

As part of the routing options configuration for Router A, you need to configure routing table groups to enable VPN route resolution in the inet.3 routing table.

Configure the routing options on Router A:

```
[edit routing-options]
interface-routes {
  rib-group inet if-rib;
}
rib inet.3 {
  static {
    route 10.255.14.178/32 next-hop gr-1/1/0.0;
  }
}
rib-groups {
  if-rib {
    import-rib [ inet.0 inet.3 ];
  }
}
```

## Configuring the Routing Options on Router D

As part of the routing options configuration for Router D, you need to configure routing table groups to enable VPN route resolution in the inet.3 routing table.

Configure the routing options on Router D:

```
[edit routing-options]
interface-routes {
  rib-group inet if-rib;
}
rib inet.3 {
  static {
    route 10.255.14.176/32 next-hop gr-1/1/0.0;
  }
}
rib-groups {
  if-rib {
    import-rib [ inet.0 inet.3 ];
  }
}
```

## Configuration Summary for Router A

<b>Configure the Routing Instance</b>	<pre>gre-config {   instance-type vrf;   interface fe-1/0/0.0;   route-distinguisher 10.255.14.176:69;   vrf-import import-config;   vrf-export export-config;   protocols {     ospf {</pre>
---------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

	<pre> export import-config; area 0.0.0.0 {     interface all; } } } </pre>
Configure MPLS	<pre> mpls {     interface all; } </pre>
Configure BGP	<pre> bgp {     traceoptions {         file bgp.trace world-readable;         flag update detail;     }     group pe-to-pe {         type internal;         neighbor 10.255.14.178 {             family inet-vpn {                 unicast;             }         }     } } </pre>
Configure OSPF	<pre> ospf {     area 0.0.0.0 {         interface all;         interface gr-1/1/0.0 {             disable;         }     } } </pre>
Configure the Tunnel Interface	<pre> interface-name {     unit 0 {         tunnel {             source 10.255.14.176;             destination 10.255.14.178;         }         family inet;         family mpls;     } } </pre>
Configure Routing Options	<pre> interface-routes {     rib-group inet if-rib; } rib inet.3 {     static {         route 10.255.14.178/32 next-hop gr-1/1/0.0;     } } </pre>

```

    }
  }
  rib-groups {
    if-rib {
      import-rib [ inet.0 inet.3 ];
    }
  }
}

```

## Configuration Summary for Router D


<b>Configure the Routing Instance</b>	<pre> gre-config {   instance-type vrf;   interface fe-1/0/1.0;   route-distinguisher 10.255.14.178:69;   vrf-import import-config;   vrf-export export-config;   protocols {     ospf {       export import-config;       area 0.0.0.0 {         interface all;       }     }   } } </pre>
<b>Configure MPLS</b>	<pre> mpls {   interface all; } </pre>
<b>Configure BGP</b>	<pre> bgp {   group pe-to-pe {     type internal;     neighbor 10.255.14.176 {       family inet-vpn {         unicast;       }     }   } } </pre>
<b>Configure OSPF</b>	<pre> ospf {   traffic-engineering;   area 0.0.0.0 {     interface all;     interface fxp0.0 {       disable;     }     interface gr-1/1/0.0 {       disable;     }   } } </pre>

<b>Configure the Tunnel Interface</b>	<pre> interface-name {   unit 0 {     tunnel {       source 10.255.14.178;       destination 10.255.14.176;     }     family inet;     family mpls;   } } </pre>
<b>Configure the Routing Options</b>	<pre> interface-routes {   rib-group inet if-rib; } rib inet.3 {   static {     route 10.255.14.176/32 next-hop gr-1/1/0.0;   } } rib-groups {   if-rib {     import-rib [ inet.0 inet.3 ];   } } </pre>

## Configuring a GRE Tunnel Interface Between a PE and CE Router

This example shows how to configure a GRE tunnel interface between a PE router and a CE router. You can use this configuration to tunnel VPN traffic across a non-MPLS core network. The network topology used in this example is shown in [Figure 70 on page 721](#).

*Figure 70: GRE Tunnel Between the CE Router and the PE Router*



For this example, complete the procedures described in the following sections:

- [Configuring the Routing Instance Without the Encapsulating Interface on page 722](#)
- [Configuring the Routing Instance with the Encapsulating Interface on page 723](#)
- [Configuring the GRE Tunnel Interface on Router CE1 on page 725](#)

## Configuring the Routing Instance Without the Encapsulating Interface

You can configure the routing instance either with or without the encapsulating interface. The following sections explain how to configure the routing instance without it:

- [Configuring the Routing Instance on Router PE1 on page 722](#)
- [Configuring the GRE Tunnel Interface on Router PE1 on page 722](#)
- [Configuring the Encapsulation Interface on Router PE1 on page 723](#)

---

### Configuring the Routing Instance on Router PE1

Configure the routing instance on Router PE1:

```
[edit routing-instances]
vpna {
  instance-type vrf;
  interface gr-1/2/0.0;
  route-distinguisher 10.255.14.174:1;
  vrf-import vpna-import;
  vrf-export vpna-export;
  protocols {
    bgp {
      group vpna {
        type external;
        peer-as 100;
        as-override;
        neighbor 10.49.2.1;
      }
    }
  }
}
```

---

### Configuring the GRE Tunnel Interface on Router PE1

Configure the GRE tunnel interface on Router PE1:

```
[edit interfaces gr-1/2/0]
unit 0 {
  tunnel {
    source 192.168.197.249;
    destination 192.168.197.250;
  }
  family inet {
    address 10.49.2.2/30;
  }
}
```

In this example, interface **t3-0/1/3** acts as the encapsulating interface for the GRE tunnel.

When you configure the **clear-dont-fragment-bit** statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than the maximum supported value, which is 9192.

For example:

```
user@host# show interfaces gr-1/2/0
unit 0 {
  clear-dont-fragment-bit;
  family inet {
    mtu 9100;
    address 10.10.1.1/32;
  }
  family mpls {
    mtu 9100;
  }
}
```

### Configuring the Encapsulation Interface on Router PE1

Configure the encapsulation interface on Router PE1:

```
[edit interfaces t3-0/1/3]
unit 0 {
  family inet {
    address 192.168.197.249/30;
  }
}
```

### Configuring the Routing Instance with the Encapsulating Interface

If the tunnel-encapsulating interface, **t3-0/1/3**, is also configured under the routing instance, then you need to specify the name of that routing instance under the interface definition. The system uses this routing instance to search for the tunnel destination address.

To configure the routing instance with the encapsulating interface, you perform the steps in the following sections:

- [Configuring the Routing Instance on Router PE1 on page 723](#)
- [Configuring the GRE Tunnel Interface on Router PE1 on page 724](#)
- [Configuring the Encapsulation Interface on Router PE1 on page 725](#)

### Configuring the Routing Instance on Router PE1

If you configure the tunnel-encapsulating interface under the routing instance, then configure the routing instance on Router PE1:

```
[edit routing-instances]
vpna {
  instance-type vrf;
  interface gr-1/2/0.0;
  interface t3-0/1/3.0;
  route-distinguisher 10.255.14.174:1;
  vrf-import vpna-import;
  vrf-export vpna-export;
  protocols {
    bgp {
      group vpna {
```

```
        type external;
        peer-as 100;
        as-override;
        neighbor 10.49.2.1;
    }
}
}
```

---

### Configuring the GRE Tunnel Interface on Router PE1

Configure the GRE tunnel interface on Router PE1:

```
[edit interfaces gr-1/2/0]
unit 0 {
    tunnel {
        source 192.168.197.249;
        destination 192.168.197.250;
        routing-instance {
            destination vpna;
        }
    }
    family inet {
        address 10.49.2.2/30;
    }
}
```

When you configure the **clear-dont-fragment-bit** statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than the maximum supported value, which is 9192.

For example:

```
user@host# show interfaces gr-1/2/0
unit 0 {
    clear-dont-fragment-bit;
    family inet {
        mtu 9100;
        address 10.10.1.1/32;
    }
    family mpls {
        mtu 9100;
    }
}
```

When you configure the **clear-dont-fragment-bit** statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than the maximum supported value, which is 9192.

For example:

```
user@host# show interfaces gr-1/2/0
unit 0 {
    clear-dont-fragment-bit;
    family inet {
        mtu 9100;
        address 10.10.1.1/32;
    }
}
```



```
}  
family mpls {  
    mtu 9100;  
}  
}
```

---

### Configuring the Encapsulation Interface on Router PE1

Configure the encapsulation interface on Router PE1:

```
[edit interfaces t3-0/1/3]  
unit 0 {  
    family inet {  
        address 192.168.197.249/30;  
    }  
}
```

### Configuring the GRE Tunnel Interface on Router CE1

Configure the GRE tunnel interface on Router CE1:

```
[edit interfaces gr-1/2/0]  
unit 0 {  
    tunnel {  
        source 192.168.197.250;  
        destination 192.168.197.249;  
    }  
    family inet {  
        address 10.49.2.1/30;  
    }  
}
```



## CHAPTER 40

# Configuring Next-Hop-Based Dynamic Tunnels for Layer 3 VPNs

- [Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels on page 727](#)
- [Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels on page 740](#)
- [Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview on page 753](#)
- [Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels on page 756](#)

### Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels

---

This example shows how to configure a dynamic generic routing encapsulation (GRE) tunnel that includes a tunnel composite next hop instead of an interface next hop. The next-hop-based dynamic GRE tunnel has a scaling advantage over the interface-based dynamic GRE tunnels.

- [Requirements on page 727](#)
- [Overview on page 728](#)
- [Configuration on page 730](#)
- [Verification on page 735](#)
- [Troubleshooting on page 739](#)

### Requirements

This example uses the following hardware and software components:

- Five MX Series routers with MPCs and MICs.
- Junos OS Release 16.2 or later running on the provider edge routers.

Before you begin:

1. Configure the device interfaces, including the loopback interface.
2. Configure the router ID and autonomous system number for the device.
3. Establish an internal BGP (IBGP) session with the remote PE device.
4. Establish OSPF peering among the devices.

## Overview

Starting with Junos OS Release 16.2, a dynamic generic routing encapsulation (GRE) tunnel supports the creation of a tunnel composite next hop for every GRE tunnel configured.

By default, for every new dynamic GRE tunnel configured, a corresponding logical tunnel interface is created. This is called an interface-based dynamic tunnel and is the default mode for creating dynamic tunnels. With the interface-based dynamic tunnel, the number of tunnels that can be configured on a device is dependant on the total number of interfaces supported on the device. The next-hop-based dynamic GRE tunnel removes the dependency on physical interfaces, and the GRE tunnel encapsulation is implemented as a next-hop instruction without creating a logical tunnel interface. This provides a scaling advantage for the number of dynamic tunnels that can be created on a device.

Starting in Junos OS Release 17.1, on MX Series routers with MPCs and MICs, the scaling limit of next-hop-based dynamic GRE tunnels is increased. The increased scaling values benefits data center networks, where a gateway router is required to communicate with a number of servers over an IP infrastructure; for example, in Contrail networking.

At a given point in time, either the next-hop-based dynamic tunnel or the default interface-based dynamic GRE tunnel can exist on a device. A switchover from one tunnel mode to another deletes the existing tunnels and creates new tunnels in the new tunnel mode according to the supported scaling limits. Similarly, at a given point in time, for the same tunnel destination, the next-hop-based tunnel encapsulation type can either be GRE or UDP.

To enable the next-hop-based dynamic GRE tunnel, include the **next-hop-based-tunnel** statement at the **[edit routing-options dynamic-tunnels gre]** hierarchy level.

The existing dynamic tunnel feature requires complete static configuration. Currently, the tunnel information received from peer devices in advertised routes is ignored. Starting in Junos OS Release 17.4R1, on MX Series routers, the next-hop-based dynamic GRE tunnels are signaled using BGP encapsulation extended community. BGP export policy is used to specify the tunnel types, advertise the sender side tunnel information, and parse and convey the receiver side tunnel information. A tunnel is created according to the received type tunnel community.

Multiple tunnel encapsulations are supported by BGP. On receiving multiple capability, the next-hop-based dynamic tunnel is created based on the configured BGP policy and tunnel preference. The tunnel preference should be consistent across both the tunnel ends for the tunnel to be set up. By default, MPLS-over-UDP (MPLSoUDP) tunnel is preferred over GRE tunnels. If dynamic tunnel configuration exists, it takes precedence over received tunnel community.

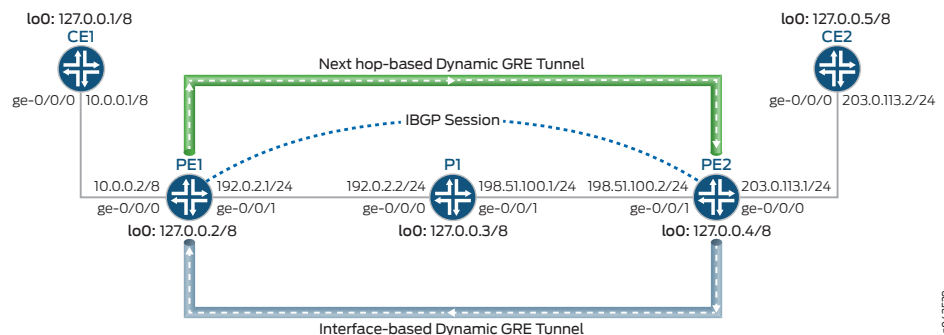
When configuring a next-hop-based dynamic GRE tunnel, be aware of the following considerations:

- Dynamic tunnels creation is triggered in the IBGP protocol next-hop resolution process.
- A switchover from the next-hop-based tunnel mode to the interface-based tunnel mode (and vice versa) is allowed, and this can impact network performance in terms of the supported IP tunnel scaling values in each mode.
- RSVP automatic mesh tunnel is not supported with the next-hop-based dynamic tunnel configuration.
- Dynamic GRE tunnel creation based upon new IPv4-mapped-IPv6 next hops is supported with this feature.
- The following features are not supported with the next-hop-based dynamic GRE tunnel configuration:
  - RSVP automatic mesh
  - Logical systems
  - Per-tunnel traffic statistics collection on the sender (MX Series) side
  - QoS enforcement, such as policing and shaping, at the tunnel interface.
  - Path maximum transmission unit discovery
  - GRE key feature
  - GRE Operation, Administration, and Maintenance (OAM) for GRE keepalive messages.

## Topology

Figure 71 on page 729 illustrates a Layer 3 VPN scenario over dynamic GRE tunnels. The customer edge (CE) devices, CE1 and CE2, connect to provider edge (PE) devices, PE1 and PE2, respectively. The PE devices are connected to a provider device (Device P1), and an internal BGP (IBGP) session interconnects the two PE devices. Two dynamic GRE tunnels are configured between the PE devices. The dynamic tunnel from Device PE1 to Device PE2 is based on the next hop tunnel mode, and the dynamic tunnel from Device PE2 to Device PE1 is based on the interface tunnel mode.

**Figure 71: Layer 3 VPN over Dynamic GRE Tunnels**



The next-hop-based dynamic GRE tunnel is handled as follows:

1. After a next-hop-based dynamic GRE tunnel is configured, a tunnel destination mask route with a tunnel composite next hop is created for the protocol next hops in the inet.3 routing table. This IP tunnel route is withdrawn only when the dynamic tunnel configuration is deleted.

The tunnel composite next hop attributes include the following:

- When Layer 3 VPN composite next hop is disabled—Source and destination address, encapsulation string, and VPN label.
  - When Layer 3 VPN composite next hop and per-prefix VPN label allocation are enabled—Source address, destination address, and encapsulation string.
  - When Layer 3 VPN composite next hop is enabled and per-prefix VPN label allocation is disabled—Source address, destination address, and encapsulation string. The route in this case is added to the other virtual routing and forwarding instance table with a secondary route.
2. The PE devices are interconnected using an IBGP session. The IBGP route next hop to a remote BGP neighbor is the protocol next hop, which is resolved using the tunnel mask route with the tunnel next hop.
  3. After the protocol next hop is resolved over the tunnel composite next hop, indirect next hops with forwarding next hops are created.
  4. The tunnel composite next hop is used to forward the next hops of the indirect next hops.

## 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.

```
CE1 set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/8
    set interfaces lo0 unit 0 family inet address 127.0.0.1/8
    set routing-options router-id 127.0.0.1
    set routing-options autonomous-system 200
    set protocols bgp group ce1-pe1 export export-loopback-direct
    set protocols bgp group ce1-pe1 peer-as 100
    set protocols bgp group ce1-pe1 neighbor 10.0.0.2
    set policy-options policy-statement export-loopback-direct term term-1 from interface
        lo0.0
    set policy-options policy-statement export-loopback-direct term term-1 from route-filter
        127.0.0.1/8 exact
    set policy-options policy-statement export-loopback-direct term term-1 then accept

CE2 set interfaces ge-0/0/0 unit 0 family inet address 203.0.113.2/24
    set interfaces lo0 unit 0 family inet address 127.0.0.5/8
```

```

set routing-options router-id 127.0.0.5
set routing-options autonomous-system 200
set protocols bgp group ce1-pe1 export export-loopback-direct
set protocols bgp group ce1-pe1 peer-as 100
set protocols bgp group ce1-pe1 neighbor 203.0.113.1
set policy-options policy-statement export-loopback-direct term term-1 from interface
  lo0.0
set policy-options policy-statement export-loopback-direct term term-1 from route-filter
  127.0.0.5/8 exact
set policy-options policy-statement export-loopback-direct term term-1 then accept

PE1 set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.2/8
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.2/8
set routing-options static route 33.0.0.0/8 next-hop 192.0.2.2
set routing-options router-id 127.0.0.2
set routing-options autonomous-system 100
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels gre-dyn-tunnel-to-pe2 source-address 127.0.0.2
set routing-options dynamic-tunnels gre-dyn-tunnel-to-pe2 gre
set routing-options dynamic-tunnels gre-dyn-tunnel-to-pe2 destination-networks
  127.0.0.0/8
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 127.0.0.2
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 127.0.0.4
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-instances L3VPN-Over-GRE-PE1 instance-type vrf
set routing-instances L3VPN-Over-GRE-PE1 interface ge-0/0/0.0
set routing-instances L3VPN-Over-GRE-PE1 route-distinguisher 127.0.0.2:1
set routing-instances L3VPN-Over-GRE-PE1 vrf-target target:600:1
set routing-instances L3VPN-Over-GRE-PE1 protocols bgp group pe1-ce1 peer-as 200
set routing-instances L3VPN-Over-GRE-PE1 protocols bgp group pe1-ce1 neighbor 10.0.0.1
  as-override

P1 set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.3/8
set routing-options router-id 127.0.0.3
set routing-options autonomous-system 100
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

PE2 set interfaces ge-0/0/0 unit 0 family inet address 203.0.113.1/24
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.2/24
set interfaces lo0 unit 0 family inet address 127.0.0.4/8
set routing-options nonstop-routing
set routing-options router-id 127.0.0.4
set routing-options autonomous-system 100

```

```

set routing-options dynamic-tunnels gre-dyn-tunnel-to-pe1 source-address 127.0.0.4
set routing-options dynamic-tunnels gre-dyn-tunnel-to-pe1 gre
set routing-options dynamic-tunnels gre-dyn-tunnel-to-pe1 destination-networks
  127.0.0.0/8
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 127.0.0.4
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 127.0.0.2
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-instances L3VPN-Over-GRE-PE2 instance-type vrf
set routing-instances L3VPN-Over-GRE-PE2 interface ge-0/0/0.0
set routing-instances L3VPN-Over-GRE-PE2 route-distinguisher 127.0.0.4:1
set routing-instances L3VPN-Over-GRE-PE2 vrf-target target:600:1
set routing-instances L3VPN-Over-GRE-PE2 protocols bgp group ebgp peer-as 200
set routing-instances L3VPN-Over-GRE-PE2 protocols bgp group ebgp neighbor 203.0.113.2
  as-override

```

**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 PE1:

1. Configure the device interfaces including the loopback interface of the device.

```

[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 10.0.0.2/8

user@PE1# set ge-0/0/1 unit 0 family inet address 192.0.2.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls

user@PE1# set lo0 unit 0 family inet address 127.0.0.2/8

```

2. Configure a static route for routes from Device PE1 with Device P1 as the next-hop destination.

```

[edit routing-options]
user@PE1# set static route 33.0.0.0/8 next-hop 192.0.2.2

```

3. Configure the router ID and autonomous system number for Device PE1.

```

[edit routing-options]
user@PE1# set router-id 127.0.0.2
user@PE1# set autonomous-system 100

```

4. Configure IBGP peering between the PE devices.

```

[edit protocols]
user@PE1# set bgp group IBGP type internal
user@PE1# set bgp group IBGP local-address 127.0.0.2
user@PE1# set bgp group IBGP family inet-vpn unicast
user@PE1# set bgp group IBGP neighbor 127.0.0.4

```



5. Configure OSPF on all the interfaces of Device PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set ospf area 0.0.0.0 interface ge-0/0/1.0
user@PE1# set ospf area 0.0.0.0 interface lo0.0 passive
```

6. Enable next-hop-based dynamic GRE tunnel configuration on Device PE1.

```
[edit routing-options]
user@PE1# set dynamic-tunnels gre next-hop-based-tunnel
```

7. Configure the dynamic GRE tunnel parameters from Device PE1 to Device PE2.

```
[edit routing-options]
user@PE1# set dynamic-tunnels gre-dyn-tunnel-to-pe2 source-address 127.0.0.2
user@PE1# set dynamic-tunnels gre-dyn-tunnel-to-pe2 gre
user@PE1# set dynamic-tunnels gre-dyn-tunnel-to-pe2 destination-networks
55.66.77.0/24
```

8. Export the load-balancing policy to the forwarding table.

```
[edit routing-options]
user@PE1# set forwarding-table export pplb
```

9. Configure a VRF routing instance on Device PE1 and other routing instance parameters.

```
[edit routing-instances]
user@PE1# set L3VPN-Over-GRE-PE1 instance-type vrf
user@PE1# set L3VPN-Over-GRE-PE1 interface ge-0/0/0.0
user@PE1# set L3VPN-Over-GRE-PE1 route-distinguisher 127.0.0.2:1
user@PE1# set L3VPN-Over-GRE-PE1 vrf-target target:600:1
```

10. Enable BGP in the routing instance configuration for peering with Device CE1.

```
[edit routing-instances]
user@PE1# set L3VPN-Over-GRE-PE1 protocols bgp group pe1-ce1 peer-as 200
user@PE1# set L3VPN-Over-GRE-PE1 protocols bgp group pe1-ce1 neighbor 10.0.0.1
as-override
```

## Results

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

```
user@PE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/8;
```

```
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.2/8;
    }
  }
}

user@PE1# show routing-options
static {
  route 33.0.0.0/8 next-hop 192.0.2.2;
}
router-id 127.0.0.2;
autonomous-system 100;
dynamic-tunnels {
  gre next-hop-based-tunnel;
  gre-dyn-tunnel-to-pe2 {
    source-address 127.0.0.2;
    gre;
    destination-networks {
      127.0.0.0/8;
    }
  }
}

user@PE1# show protocols
bgp {
  group IBGP {
    type internal;
    local-address 127.0.0.2;
    family inet-vpn {
      unicast;
    }
    neighbor 127.0.0.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}

user@PE1# show routing-instances
```

```
L3VPN-Over-GRE-PE1 {  
  instance-type vrf;  
  interface ge-0/0/0.0;  
  route-distinguisher 127.0.0.2:1;  
  vrf-target target:600:1;  
  protocols {  
    bgp {  
      group pe1-cel {  
        peer-as 200;  
        neighbor 10.0.0.1 {  
          as-override;  
        }  
      }  
    }  
  }  
}
```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between PE Devices on page 735](#)
- [Verify the Dynamic Tunnel Routes on Device PE1 on page 736](#)
- [Verify the Dynamic Tunnel Routes on Device PE2 on page 737](#)
- [Verifying That the Routes Have the Expected Indirect-Next-Hop Flag on page 738](#)

### Verifying the Connection Between PE Devices

**Purpose** Verify the BGP peering status between Device PE1 and Device PE2, and the BGP routes received from Device PE2.

**Action** From operational mode, run the **show bgp summary** and **show route receive-protocol bgp ip-address table bgp.l3vpn.0** commands.

```

user@PE1> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State    Pending
bgp.l3vpn.0
                2          2          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
127.0.0.4      100      139      136       0       0      58:23 Establ

    bgp.l3vpn.0: 2/2/2/0
    L3VPN-Over-GRE-PE1.inet.0: 2/2/2/0
10.0.0.1       200      135      136       0       0      58:53
Establ
    L3VPN-Over-GRE-PE1.inet.0: 1/1/1/0

user@PE1> show route receive-protocol bgp 127.0.0.4 table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref    AS path
127.0.0.4:1:127.0.0.5/8
*                127.0.0.4          100      200 I
127.0.0.4:1:203.0.113.0/24
*                127.0.0.4          100      I

```

- Meaning**
- In the first output, the BGP session state is **Establ**, which means that the session is up and the PE devices are peered.
  - In the second output, Device PE1 has learned two BGP routes from Device PE2.

#### Verify the Dynamic Tunnel Routes on Device PE1

- Purpose** Verify the routes in the inet.3 routing table and the dynamic tunnel database information on Device PE1.

**Action** From operational mode, run the **show route table inet.3** and **show dynamic-tunnels database terse** commands.

```

user@PE1> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.0/8      *[Tunnel/300] 01:01:45
                  Tunnel
127.0.0.4/8      *[Tunnel/300] 00:10:24
                  Tunnel Composite

user@PE1> show dynamic-tunnels database terse
Table: inet.3

Destination-network: 127.0.0.0/8
Destination      Source      Next-hop      Type      Status
127.0.0.4/8      127.0.0.2  0xb395e70 nhid 612  gre      Up

```

- Meaning**
- In the first output, because Device PE1 is configured with the next-hop-based dynamic GRE tunnel, a tunnel composite route is created for the inet.3 routing table route entry.
  - In the second output, the dynamic GRE tunnel created from Device PE1 to Device PE2 is up and has a next-hop ID assigned to it instead of a next-hop interface.

#### Verify the Dynamic Tunnel Routes on Device PE2

**Purpose** Verify the routes in the inet.3 routing table and the dynamic tunnel database information on Device PE2.

**Action** From operational mode, run the **show route table inet.3** and **show dynamic-tunnels database terse** commands.

```

user@PE2> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.0/8      *[Tunnel/300] 01:06:52
                  Tunnel
127.0.0.2/8      *[Tunnel/300] 01:04:45
                  > via gr-0/1/0.32769

user@PE1> show dynamic-tunnels database terse
Table: inet.3

Destination-network: 127.0.0.0/8
Destination      Source      Next-hop      Type      Status
127.0.0.2/8      127.0.0.4   gr-0/1/0.32769 gre        Up

```

- Meaning**
- In the first output, because Device PE2 has the default interface-based dynamic GRE tunnel configuration, a new interface is created for the inet.3 routing table route entry.
  - In the second output, the dynamic GRE tunnel created from Device PE2 to Device PE1 is up, and has the newly created interface name assigned to it.

### Verifying That the Routes Have the Expected Indirect-Next-Hop Flag

**Purpose** Verify that Device PE1 and Device PE2 are configured to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table.

**Action** From operational mode, run the **show krt indirect-next-hop** command on Device PE1 and Device PE2.

```

user@PE1> show krt indirect-next-hop
Indirect Nexthop:
Index: 1048574 Protocol next-hop address: 127.0.0.4
RIB Table: bgp.13vpn.0
Label: Push 299792
Policy Version: 1                      References: 2
Locks: 3                               0xb2ab630
Flags: 0x0
INH Session ID: 0x0
INH Version ID: 0
Ref RIB Table: unknown
  Tunnel type: GRE, Reference count: 3, nhid: 612
  Destination address: 127.0.0.4, Source address: 127.0.0.2
  Tunnel id: 1, VPN Label: Push 299792, TTL action: prop-ttl
IGP FRR Interesting proto count : 2
Chain IGP FRR Node Num          : 1
  IGP Resolver node(hex)         : 0xb3c6d9c
  IGP Route handle(hex)          : 0xb1ad230      IGP rt_entry protocol
: Tunnel

```

```

IGP Actual Route handle(hex) : 0x0          IGP Actual rt_entry protocol
: Any

user@PE2> show krt indirect-next-hop
Indirect Nexthop:
Index: 1048574 Protocol next-hop address: 127.0.0.2
RIB Table: bgp.13vpn.0
Label: Push 299792
Policy Version: 2                      References: 2
Locks: 3                               0xb2ab630
Flags: 0x2
INH Session ID: 0x145
INH Version ID: 0
Ref RIB Table: unknown
  Next hop: via gr-0/1/0.32769
  Label operation: Push 299792
  Label TTL action: prop-ttl
  Load balance label: Label 299792: None;
  Label element ptr: 0xb395d40
  Label parent element ptr: 0x0
  Label element references: 1
  Label element child references: 0
  Label element lsp id: 0
  Session Id: 0x144
  IGP FRR Interesting proto count : 2
  Chain IGP FRR Node Num          : 1
    IGP Resolver node(hex)        : 0xb3d36e8
    IGP Route handle(hex)         : 0xb1af060    IGP rt_entry protocol
: Tunnel
  IGP Actual Route handle(hex) : 0x0          IGP Actual rt_entry protocol
: Any

```

- Meaning**
- In the first output, Device PE1 has a next-hop-based dynamic GRE tunnel to Device PE2.
  - In the second output, Device PE2 has an interface-based dynamic GRE tunnel to device PE1.

## Troubleshooting

To troubleshoot the next-hop-based dynamic tunnels, see:

- [Troubleshooting Commands on page 739](#)

### Troubleshooting Commands

**Problem** The next-hop-based dynamic GRE tunnel configuration is not taking effect.

**Solution** To troubleshoot the next-hop-based GRE tunnel configuration, use the following **traceroute** commands at the **[edit routing-options dynamic-tunnels]** statement hierarchy:

- **traceoptions file *file-name***
- **traceoptions file size *file-size***

- `traceoptions flag all`

For example:

```
[edit routing-options dynamic-tunnels]
traceoptions {
  file gre_dyn_pe1.wri size 4294967295;
  flag all;
}
```

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, on MX Series routers, the next-hop-based dynamic GRE tunnels are signaled using BGP encapsulation extended community.
17.1R1	Starting in Junos OS Release 17.1, on MX Series routers with MPCs and MICs, the scaling limit of next-hop-based dynamic GRE tunnels is increased.

**Related  
Documentation**

- [Configuring GRE Tunnels for Layer 3 VPNs on page 711](#)
- [Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels on page 740](#)

## Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels

This example shows how to configure a dynamic MPLS-over-UDP tunnel that includes a tunnel composite next hop. The MPLS-over-UDP feature provides a scaling advantage on the number of IP tunnels supported on a device.

- [Requirements on page 740](#)
- [Overview on page 741](#)
- [Configuration on page 743](#)
- [Verification on page 748](#)
- [Troubleshooting on page 752](#)

### Requirements

This example uses the following hardware and software components:

- Five MX Series routers with MPCs and MICs.
- Junos OS Release 16.2 or later running on the PE routers.

Before you begin:

1. Configure the device interfaces, including the loopback interface.
2. Configure the router ID and autonomous system number for the device.



3. Establish an internal BGP (IBGP) session with the remote PE device.
4. Establish OSPF peering among the devices.

## Overview

Starting with Junos OS Release 16.2, a dynamic UDP tunnel supports the creation of a tunnel composite next hop for every UDP tunnel configured. These next-hop-based dynamic UDP tunnels are referred to as MPLS-over-UDP tunnels. The tunnel composite next hop are enabled by default for the MPLS-over-UDP tunnels.

Starting in Junos OS Release 17.1, on MX Series routers with MPCs and MICs, the scaling limit of MPLS-over-UDP tunnels is increased.



**NOTE:** For the same tunnel destination, the next-hop-based dynamic tunnel encapsulation can either be GRE or UDP. Having both tunnel encapsulations for the same tunnel destination causes a commit error under the dynamic-tunnel configuration.

The existing dynamic tunnel feature requires complete static configuration. Currently, the tunnel information received from peer devices in advertised routes is ignored. Starting in Junos OS Release 17.4R1, on MX Series routers, the next-hop-based dynamic MPLS-over-UDP tunnels are signaled using BGP encapsulation extended community. BGP export policy is used to specify the tunnel types, advertise the sender side tunnel information, and parse and convey the receiver side tunnel information. A tunnel is created according to the received type tunnel community.

Multiple tunnel encapsulations are supported by BGP. On receiving multiple capability, the next-hop-based dynamic tunnel is created based on the configured BGP policy and tunnel preference. The tunnel preference should be consistent across both the tunnel ends for the tunnel to be set up. By default, MPLS-over-UDP (MPLSoUDP) tunnel is preferred over GRE tunnels. If dynamic tunnel configuration exists, it takes precedence over received tunnel community.

When configuring a next-hop-based dynamic MPLS-over-UDP tunnel, be aware of the following considerations:

- An IBGP session must be configured between the PE devices.
- A switchover between the next-hop-based dynamic tunnel encapsulations (UDP and GRE) is allowed, and this can impact network performance in terms of the supported IP tunnel scaling values in each mode.
- Having both GRE and UDP next-hop-based dynamic tunnel encapsulation types for the same tunnel destination leads to a commit failure.
- Graceful Routing Engine switchover (GRES) is supported with MPLS-over-UDP, and the MPLS-over-UDP tunnel type flags are unified ISSU and NSR compliant.
- MPLS-over-UDP tunnels are supported on virtual MX (vMX).

- MPLS-over-UDP tunnels support dynamic GRE tunnel creation based upon new IPv4-mapped-IPv6 next hops.
- MPLS-over-UDP tunnels are supported in interoperability with contrail, wherein the MPLS-over-UDP tunnels are created from the contrail vRouter to an MX gateway. To enable this, the following community is required to be advertised in the route from the MX Series router to the contrail vRouter:

```
[edit policy-options community]
  udp members 0x030c:64512:13;
```


At a given point in time, only one tunnel type is supported on the contrail vRouter—next-hop-based dynamic GRE tunnels, MPLS-over-UDP tunnels, or VXLAN.

- The following features are not supported with the next-hop-based dynamic MPLS-over-UDP tunnel configuration:
  - RSVP automatic mesh
  - Plain IPV6 GRE and UDP tunnel configuration
  - Logical systems

## Topology

Figure 72 on page 742 illustrates a Layer 3 VPN scenario over dynamic MPLS-over-UDP tunnels. The customer edge (CE) devices, CE1 and CE2, connect to provider edge (PE) devices, PE1 and PE2, respectively. The PE devices are connected to a provider device (Device P1), and an internal BGP (IBGP) session interconnects the two PE devices. Two dynamic next-hop-based MPL-over-UDP tunnels are configured between the PE devices.

**Figure 72: Dynamic MPLS-over-UDP Tunnels**



The MPLS-over-UDP tunnel is handled as follows:

1. After a MPLS-over-UDP tunnel is configured, a tunnel destination mask route with a tunnel composite next hop is created for the tunnel in the inet.3 routing table. This IP tunnel route is withdrawn only when the dynamic tunnel configuration is deleted.

The tunnel composite next-hop attributes include the following:

- When Layer 3 VPN composite next hop is disabled—Source and destination address, encapsulation string, and VPN label.
  - When Layer 3 VPN composite next hop and per-prefix VPN label allocation are enabled—Source address, destination address, and encapsulation string.
  - When Layer 3 VPN composite next hop is enabled and per-prefix VPN label allocation is disabled—Source address, destination address, and encapsulation string. The route in this case is added to the other virtual routing and forwarding instance table with a secondary route.
2. The PE devices are interconnected using an IBGP session. The IBGP route next hop to a remote BGP neighbor is the protocol next hop, which is resolved using the tunnel mask route with the tunnel next hop.
  3. After the protocol next hop is resolved over the tunnel composite next hop, indirect next hops with forwarding next hops are created.
  4. The tunnel composite next hop is used to forward the next hops of the indirect next hops.

## 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.

```

CE1  set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/8
      set interfaces lo0 unit 0 family inet address 127.0.0.1/8
      set routing-options router-id 127.0.0.1
      set routing-options autonomous-system 200
      set protocols bgp group ce1-pe1 export export-loopback-direct
      set protocols bgp group ce1-pe1 peer-as 100
      set protocols bgp group ce1-pe1 neighbor 10.0.0.2
      set policy-options policy-statement export-loopback-direct term term-1 from interface
        lo0.0
      set policy-options policy-statement export-loopback-direct term term-1 from route-filter
        127.0.0.1/8 exact
      set policy-options policy-statement export-loopback-direct term term-1 then accept

CE2  set interfaces ge-0/0/0 unit 0 family inet address 203.0.113.2/24
      set interfaces lo0 unit 0 family inet address 127.0.0.5/8
      set routing-options router-id 127.0.0.5

```

```

set routing-options autonomous-system 200
set protocols bgp group ce1-pe1 export export-loopback-direct
set protocols bgp group ce1-pe1 peer-as 100
set protocols bgp group ce1-pe1 neighbor 203.0.113.1
set policy-options policy-statement export-loopback-direct term term-1 from interface
  lo0.0
set policy-options policy-statement export-loopback-direct term term-1 from route-filter
  127.0.0.5/8 exact
set policy-options policy-statement export-loopback-direct term term-1 then accept

```

```

PE1 set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.2/8
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.2/8
set routing-options static route 33.0.0.0/8 next-hop 192.0.2.2
set routing-options router-id 127.0.0.2
set routing-options autonomous-system 100
set routing-options forwarding-table export pplib
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe2 source-address 127.0.0.2
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe2 udp
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe2 destination-networks
  127.0.0.0/8
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 127.0.0.2
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 127.0.0.4
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-instances MPLS-over-UDP-PE1 instance-type vrf
set routing-instances MPLS-over-UDP-PE1 interface ge-0/0/0.0
set routing-instances MPLS-over-UDP-PE1 route-distinguisher 127.0.0.2:1
set routing-instances MPLS-over-UDP-PE1 vrf-target target:600:1
set routing-instances MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 peer-as 200
set routing-instances MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 neighbor 10.0.0.1
  as-override

```

```

P1 set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.3/8
set routing-options router-id 127.0.0.3
set routing-options autonomous-system 100
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

```

PE2 set interfaces ge-0/0/0 unit 0 family inet address 203.0.113.1/24
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.2/24
set interfaces lo0 unit 0 family inet address 127.0.0.4/8
set routing-options nonstop-routing
set routing-options router-id 127.0.0.4
set routing-options autonomous-system 100

```

```

set routing-options forwarding-table export pplb
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe1 source-address 127.0.0.4
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe1 udp
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe1 destination-networks
  127.0.0.0/8
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 127.0.0.4
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 127.0.0.2
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-instances MPLS-over-UDP-PE2 instance-type vrf
set routing-instances MPLS-over-UDP-PE2 interface ge-0/0/0.0
set routing-instances MPLS-over-UDP-PE2 route-distinguisher 127.0.0.4:1
set routing-instances MPLS-over-UDP-PE2 vrf-target target:600:1
set routing-instances MPLS-over-UDP-PE2 protocols bgp group ebgp peer-as 200
set routing-instances MPLS-over-UDP-PE2 protocols bgp group ebgp neighbor 203.0.113.2
  as-override

```

**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 PE1:

1. Configure the device interfaces including the loopback interface of the device.

```

[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 10.0.0.2/8

user@PE1# set ge-0/0/1 unit 0 family inet address 192.0.2.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls

user@PE1# set lo0 unit 0 family inet address 127.0.0.2/8

```

2. Configure a static route for routes from Device PE1 with Device P1 as the next-hop destination.

```

[edit routing-options]
user@PE1# set static route 33.0.0.0/8 next-hop 192.0.2.2

```

3. Configure the router-ID and autonomous system number for Device PE1.

```

[edit routing-options]
user@PE1# set router-id 127.0.0.2
user@PE1# set autonomous-system 100

```

4. Configure IBGP peering between the PE devices.

```

[edit protocols]
user@PE1# set bgp group IBGP type internal
user@PE1# set bgp group IBGP local-address 127.0.0.2
user@PE1# set bgp group IBGP family inet-vpn unicast
user@PE1# set bgp group IBGP neighbor 127.0.0.4

```

5. Configure OSPF on all the interfaces of Device PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set ospf area 0.0.0.0 interface ge-0/0/1.0
user@PE1# set ospf area 0.0.0.0 interface lo0.0 passive
```

6. Enable next-hop-based dynamic GRE tunnel configuration on Device PE1.



**NOTE:** This step is required only for illustrating the implementation difference between next-hop-based dynamic GRE tunnels and MPLS-over-UDP tunnels.

```
[edit routing-options]
user@PE1# set dynamic-tunnels gre next-hop-based-tunnel
```

7. Configure the MPLS-over-UDP tunnel parameters from Device PE1 to Device PE2.

```
[edit routing-options]
user@PE1# set dynamic-tunnels udp-dyn-tunnel-to-pe2 source-address 127.0.0.2
user@PE1# set dynamic-tunnels udp-dyn-tunnel-to-pe2 udp
user@PE1# set dynamic-tunnels udp-dyn-tunnel-to-pe2 destination-networks
127.0.0.0/8
```

8. Configure a VRF routing instance on Device PE1 and other routing instance parameters.

```
[edit routing-instances]
user@PE1# set MPLS-over-UDP-PE1 instance-type vrf
user@PE1# set MPLS-over-UDP-PE1 interface ge-0/0/0.0
user@PE1# set MPLS-over-UDP-PE1 route-distinguisher 127.0.0.2:1
user@PE1# set MPLS-over-UDP-PE1 vrf-target target:600:1
```

9. Enable BGP in the routing instance configuration for peering with Device CE1.

```
[edit routing-instances]
user@PE1# set MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 peer-as 200
user@PE1# set MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 neighbor 10.0.0.1
as-override
```

---

## Results

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

```
user@PE1# show interfaces
ge-0/0/0 {
```

```
unit 0 {
  family inet {
    address 10.0.0.2/8;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.2/8;
    }
  }
}

user@PE1# show routing-options
static {
  route 33.0.0.0/8 next-hop 192.0.2.2;
}
router-id 127.0.0.2;
autonomous-system 100;
forwarding-table {
  export pplb;
}
dynamic-tunnels {
  gre next-hop-based-tunnel;
  udp-dyn-tunnel-to-pe2 {
    source-address 127.0.0.2;
    udp;
    destination-networks {
      127.0.0.0/8;
    }
  }
}

user@PE1# show protocols
bgp {
  group IBGP {
    type internal;
    local-address 127.0.0.2;
    family inet-vpn {
      unicast;
    }
    neighbor 127.0.0.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface lo0.0 {
```

```
        passive;
    }
}

user@PE1# show routing-instances
MPLS-over-UDP-PE1 {
  instance-type vrf;
  interface ge-0/0/0.0;
  route-distinguisher 127.0.0.2:1;
  vrf-target target:600:1;
  protocols {
    bgp {
      group pe1-ce1 {
        peer-as 200;
        neighbor 10.0.0.1 {
          as-override;
        }
      }
    }
  }
}
```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between PE Devices on page 748](#)
- [Verify the Dynamic Tunnel Routes on Device PE1 on page 749](#)
- [Verify the Dynamic Tunnel Routes on Device PE2 on page 751](#)
- [Verifying That the Routes Have the Expected Indirect-Next-Hop Flag on page 751](#)

### Verifying the Connection Between PE Devices

**Purpose** Verify the BGP peering status between Device PE1 and Device PE2, and the BGP routes received from Device PE2.



**Action** From operational mode, run the **show bgp summary** and **show route receive-protocol bgp ip-address table bgp.l3vpn.0** commands.

```

user@PE1> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State    Pending
bgp.l3vpn.0
                2          2          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
127.0.0.4      100      139      136       0       0      58:23 Establ

    bgp.l3vpn.0: 2/2/2/0
    MPLS-over-UDP-PE1.inet.0: 2/2/2/0
10.0.0.1       200      135      136       0       0      58:53
Establ
    MPLS-over-UDP-PE1.inet.0: 1/1/1/0

user@PE1> show route receive-protocol bgp 127.0.0.4 table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop      MED    Lclpref    AS path
127.0.0.4:1:127.0.0.5/8
*                127.0.0.4          100      200 I
127.0.0.4:1:200.1.1.0/24
*                127.0.0.4          100      I

```

- Meaning**
- In the first output, the BGP session state is **Establ**, which means that the session is up and the PE devices are peered.
  - In the second output, Device PE1 has learned two BGP routes from Device PE2.

#### Verify the Dynamic Tunnel Routes on Device PE1

**Purpose** Verify the routes in the inet.3 routing table and the dynamic tunnel database information on Device PE1.

**Action** From operational mode, run the **show route table inet.3**, **show dynamic-tunnels database terse**, **show dynamic-tunnels database**, and **show dynamic-tunnels database summary** commands.

```

user@PE1> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.0/8      *[Tunnel/300] 00:21:18
                  Tunnel
127.0.0.4/8      *[Tunnel/300] 00:21:18
                  Tunnel Composite

user@PE1> show dynamic-tunnels database terse
Table: inet.3

Destination-network: 127.0.0.0/8
Destination      Source      Next-hop      Type      Status
127.0.0.4/8      127.0.0.2   0xb395b10 nhid 613   udp      Up

user@PE1> show dynamic-tunnels database
Table: inet.3

Destination-network: 55.0.0.0/8

Destination-network: 55.66.0.0/16

Destination-network: 55.66.77.0/24
Tunnel to: 127.0.0.4/8
Reference count: 2
Next-hop type: UDP
Source address: 127.0.0.2 Tunnel Id: 2
Next hop: tunnel-composite, 0xb395b10, nhid 613
VPN Label: Push 299776 Reference count: 3
Traffic Statistics: Packets 0, Bytes 0
State: Up

user@PE1> show dynamic-tunnels database summary
Dynamic Tunnels, Total 1 displayed
GRE Tunnel:
Active Tunnel Mode, Next Hop Base
  IFL Based, Total 0 displayed, Up 0, Down 0
  Nexthop Based, Total 0 displayed, Up 0, Down 0
RSVP Tunnel:
  Total 0 displayed
UDP Tunnel:
  Total 1 displayed, Up 1, Down 0

```

**Meaning**

- In the first output, because Device PE1 is configured with the MPLS-over-UDP tunnel, a tunnel composite route is created for the inet.3 routing table route entry.
- In the remaining outputs, the MPLS-over-UDP tunnel is displayed with the tunnel encapsulation type, tunnel next hop parameters, and tunnel status.

### Verify the Dynamic Tunnel Routes on Device PE2

**Purpose** Verify the routes in the inet.3 routing table and the dynamic tunnel database information on Device PE2.

**Action** From operational mode, run the **show route table inet.3**, and the **show dynamic-tunnels database terse** commands.

```
user@PE2> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.0/8      *[Tunnel/300] 00:39:31
                  Tunnel
127.0.0.2/8      *[Tunnel/300] 00:24:53
                  Tunnel Composite

user@PE1> show dynamic-tunnels database terse
Table: inet.3

Destination-network: 127.0.0.0/8
Destination      Source      Next-hop      Type      Status
127.0.0.2/8      127.0.0.4   0xb395450 nhid 615   udp      Up
```

**Meaning** The outputs show the MPLS-over-UDP tunnel creation and the next-hop ID assigned as the next-hop interface, similar to Device PE1.

### Verifying That the Routes Have the Expected Indirect-Next-Hop Flag

**Purpose** Verify that Device PE1 and Device PE2 are configured to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table.

**Action** From operational mode, run the **show krt indirect-next-hop** command on Device PE1 and Device PE2.

```
user@PE1> show krt indirect-next-hop
Indirect Nexthop:
Index: 1048574 Protocol next-hop address: 127.0.0.4
RIB Table: bgp.13vpn.0
Label: Push 299776
Policy Version: 1                      References: 1
Locks: 3                               0xb2ab630
Flags: 0x0
INH Session ID: 0x0
INH Version ID: 0
Ref RIB Table: unknown
  Tunnel type: UDP, Reference count: 3, nhid: 613
  Destination address: 127.0.0.4, Source address: 127.0.0.2
  Tunnel id: 2, VPN Label: Push 299776, TTL action: prop-ttl
  IGP FRR Interesting proto count : 1
  Chain IGP FRR Node Num          : 1
```

```

        IGP Resolver node(hex)      : 0xb3c70dc
        IGP Route handle(hex)       : 0xb1ae688      IGP rt_entry protocol
: Tunnel
        IGP Actual Route handle(hex) : 0x0          IGP Actual rt_entry protocol
: Any

```

```

user@PE2> show krt indirect-next-hop
Indirect Nexthop:
Index: 1048575 Protocol next-hop address: 127.0.0.2
RIB Table: bgp.l3vpn.0
Label: Push 299776
Policy Version: 1                      References: 2
Locks: 3                               0xb2ab740
Flags: 0x0
INH Session ID: 0x0
INH Version ID: 0
Ref RIB Table: unknown
    Tunnel type: UDP, Reference count: 3, nhid: 615
    Destination address: 127.0.0.2, Source address: 127.0.0.4
    Tunnel id: 1, VPN Label: Push 299776, TTL action: prop-ttl
    IGP FRR Interesting proto count : 2
    Chain IGP FRR Node Num          : 1
        IGP Resolver node(hex)      : 0xb3d3a28
        IGP Route handle(hex)       : 0xb1ae634      IGP rt_entry protocol
: Tunnel
        IGP Actual Route handle(hex) : 0x0          IGP Actual rt_entry protocol
: Any

```

**Meaning** The outputs show that a next-hop-based dynamic MPLS-over-UDP tunnel is created between the PE devices.

## Troubleshooting

To troubleshoot the next-hop-based dynamic tunnels, see:

- [Troubleshooting Commands on page 752](#)

### Troubleshooting Commands

**Problem** The next-hop-based dynamic MPLS-over-UDP tunnel configuration is not taking effect.

**Solution** To troubleshoot the next-hop-based MPLS-over-UDP tunnel configuration, use the following **traceroute** commands at the **[edit routing-options dynamic-tunnels]** statement hierarchy:

- **traceoptions file *file-name***
- **traceoptions file size *file-size***
- **traceoptions flag all**

For example:

```
[edit routing-options dynamic-tunnels]
```

```

traceoptions {
  file udp_dyn_pe1.wri size 4294967295;
  flag all;
}

```

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, on MX Series routers, the next-hop-based dynamic MPLS-over-UDP tunnels are signaled using BGP encapsulation extended community.
17.1R1	Starting in Junos OS Release 17.1, on MX Series routers with MPCs and MICs, the scaling limit of MPLS-over-UDP tunnels is increased.

**Related  
Documentation**

- [Configuring GRE Tunnels for Layer 3 VPNs on page 711](#)
- [Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels on page 727](#)

## Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview


With the rise in deployment of high-scale IP tunnels in data centers, there is a need to add security measures that allow users to limit malicious traffic from compromised virtual machines (VMs). One possible attack is the injecting of traffic into an arbitrary customer VPN from a compromised server through the gateway router. In such cases, anti-spoofing checks on IP tunnels ensure that only legitimate sources are injecting traffic into data centers from their designated IP tunnels.

Next-hop-based dynamic IP tunnels create a tunnel composite next hop for every dynamic tunnel created on the device. Because next-hop-based dynamic tunnels remove the dependency on physical interfaces for every new dynamic tunnel configured, configuring next-hop-based dynamic tunnels provides a scaling advantage over the number of dynamic tunnels that can be created on a device. Starting in Junos OS Release 17.1, anti-spoofing capabilities for next-hop-based dynamic IP tunnels is provided for next-hop-based dynamic tunnels. With this enhancement, a security measure is implemented to prevent injecting of traffic into an arbitrary customer VPN from a compromised server through the gateway router.

Anti-spoofing is implemented using reverse path forwarding checks in the Packet Forwarding Engine. The checks are implemented for the traffic coming through the tunnel to the routing instance. Currently, when the gateway router receives traffic from a tunnel, only the destination lookup is done and the packet is forwarded accordingly. When anti-spoofing protection is enabled, the gateway router also does a source address lookup of the encapsulation packet IP header in the VPN, in addition to the tunnel destination lookup. This ensures that legitimate sources are injecting traffic through their designated IP tunnels. As a result, anti-spoofing protection ensures that the tunnel traffic is received from a legitimate source on the designated tunnels.

[Figure 73 on page 754](#) illustrates a sample topology with the requirements for anti-spoofing protection.

Figure 73: Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels



In this example, the gateway router is Router G. Router G has two VPNs—Green and Blue. The two servers, Server A and Server B, can reach the Green and Blue VPNs on Router G through the next-hop-based dynamic tunnels T1 and T2, respectively. Several hosts and virtual machines (P, Q, R, S, and T) connected to the servers can reach the VPNs through the gateway router, Router G. Router G has the virtual routing and forwarding (VRF) tables for Green and Blue VPNs, each populated with the reachability information for the virtual machines in those VPNs.

For example, in VPN Green, Router G uses tunnel T1 to reach host P, tunnel T2 to reach hosts R and S, and load balancing is done between tunnels T1 and T2 to reach the multihomed host Q. In VPN Blue, Router G uses tunnel T1 to reach hosts P and R, and tunnel T2 to reach hosts Q and T.

The check passes for reverse path forwarding when:

- A packet comes from a legitimate source on its designated tunnel.  
Host P in VPN Green sends a packet to host X using tunnel T1. Because Router G can reach host P through tunnel T1, it allows the packet to pass and forwards the packet to host X.
- A packet comes from a multihomed source on its designated tunnels.  
Host Q in VPN Green is multihomed on servers A and B, and can reach Router G through tunnels T1 and T2. Host Q sends a packet to host Y using tunnel T1, and a packet to host X using tunnel T2. Because Router G can reach host Q through tunnels T1 and T2, it allows the packets to pass and forwards them to hosts Y and X, respectively.

Layer 3 VPNs do not have anti-spoofing protection enabled by default. To enable anti-spoofing for next-hop-based dynamic tunnels, include the **ip-tunnel-rpf-check** statement at the **[edit routing-instances routing-instance-name routing-options forwarding-table]** hierarchy level. The reverse path forwarding check is applied to the

VRF routing instance only. The default mode is set to **strict**, where the packet that comes from a source on a nondesignated tunnel does not pass the check. The **ip-tunnel-rpf-check** mode can be set as **loose**, where the reverse path forwarding check fails when the packet comes from a nonexistent source. An optional firewall filter can be configured under the **ip-tunnel-rpf-check** statement to count and log the packets that failed the reverse path forwarding check.

The following sample output shows an anti-spoofing configuration:

```
[edit routing-instances routing-instance-name routing-options forwarding-table]
ip-tunnel-rpf-check {
  mode loose;
  fail-filter filter-name;
}
```

Take the following guidelines under consideration when configuring anti-spoofing protection for next-hop-based dynamic tunnels:

- Anti-spoofing protection can be enabled for IPv4 tunnels and IPv4 data traffic only. The anti-spoofing capabilities are not supported on IPv6 tunnels and IPv6 data traffic.
- Anti-spoofing for next-hop-based dynamic tunnels can detect and prevent a compromised virtual machine (inner source reverse path forwarding check) but not a compromised server that is label-spoofing.
- The next-hop-based IP tunnels can originate and terminate on an inet.0 routing table.
- Anti-spoofing protection is effective when the VRF routing instance has label-switched interfaces (LSIs) (using the **vrf-table-label**), or virtual tunnel (VT) interfaces. With **per-next-hop** label on the VRF routing instance, anti-spoofing protection is not supported.
- The **rpf fail-filter** is applicable only to the inner IP packet.
- Enabling anti-spoofing checks does not affect the scaling limit of the next-hop-based dynamic tunnels on a device.
- The system resource utilization with anti-spoofing protection enabled for the VRF routing instance is slightly higher than the utilization of next-hop-based dynamic tunnels without the anti-spoofing protection enabled.
- Anti-spoofing protection requires additional source IP address checks, which has minimal impact on network performance.
- Graceful Routing Engine switchover (GRES) and in-service software upgrade (ISSU) are supported with anti-spoofing protection.

#### Related Documentation

- [ip-tunnel-rpf-check on page 1080](#)
- [Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels on page 756](#)

## Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels

This example shows how to configure reverse path forwarding checks for the virtual routing and forwarding (VRF) routing instance to enable anti-spoofing protection for next-hop-based dynamic tunnels. The checks ensure that legitimate sources are injecting traffic through their designated IP tunnels.

- [Requirements on page 756](#)
- [Overview on page 756](#)
- [Configuration on page 758](#)
- [Verification on page 763](#)

### Requirements

This example uses the following hardware and software components:

- Three MX Series Routers with MICs, each connected to a host device.
- Junos OS Release 17.1 or later running on one or all the routers.

Before you begin:

- Enable tunnel services configuration on the Flexible PIC Concentrator.
- Configure the router interfaces.
- Configure the router-ID and assign an autonomous system number for the router.
- Establish an internal BGP (IBGP) session with the tunnel endpoints.
- Configure RSVP on all the routers.
- Configure OSPF or any other interior gateway protocol on all the routers.
- Configure two dynamic next-hop-based IP tunnels between the two routers.
- Configure a VRF routing instance for every router-to-host connection.

### Overview

Starting in Junos OS Release 17.1, anti-spoofing capabilities are added to next-hop-based dynamic IP tunnels, where checks are implemented for the traffic coming through the tunnel to the routing instance using reverse path forwarding in the Packet Forwarding Engine.

Currently, when the gateway router receives traffic from a tunnel, only the destination address lookup is done before forwarding. With anti-spoofing protection, the gateway router does a source address lookup of the encapsulation packet IP header in the VPN to ensure that legitimate sources are injecting traffic through their designated IP tunnels. This is called the strict mode and is the default behavior of anti-spoofing protection. To pass traffic from nondesignated tunnels, the reverse path forwarding check is enabled in the loose mode. For traffic received from nonexistent sources, the reverse path forwarding check fails for both the strict and loose modes.




Anti-spoofing is supported on VRF routing instances. To enable anti-spoofing for dynamic tunnels, include the `ip-tunnel-rpf-check` statement at the `[edit routing-instances routing-instance-name routing-options forwarding-table]` hierarchy level.

### Topology

Figure 74 on page 757 illustrates a sample network topology enabled with anti-spoofing protection. Routers R0, R1 and R2 are each connected to hosts Host0, Host1, and Host2, respectively. Two generic routing encapsulation (GRE) next-hop-based dynamic tunnels, Tunnel 1 and Tunnel 2 – connect Router R0 with Routers R1 and R2, respectively. The VRF routing instance is running between each router and its connected host devices.

Figure 74: Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels



Taking as an example, three packets (Packets A, B, and C) are received on Router 0 from Router R2 through the next-hop-based dynamic GRE tunnel (Tunnel 2). The source IP address of these packets are 172.17.0.2 (Packet A), 172.18.0.2 (Packet B), and 172.20.0.2 (Packet C).

The source IP address of Packets A and B belong to Host 2 and Host 1, respectively. Packet C is a nonexistent source tunnel. The designated tunnel in this example is Tunnel 2, and the nondesignated tunnel is Tunnel 1. Therefore, the packets are processed as follows:

- **Packet A**—Because the source is coming from a designated tunnel (Tunnel 2), Packet A passes the reverse path forwarding check and is processed for forwarding through Tunnel 2.
- **Packet B**—Because the source is coming from Tunnel 1, which is a nondesignated tunnel, by default, Packet B fails the reverse path forwarding check in the strict mode. If loose mode is enabled, Packet B is allowed for forwarding.
- **Packet C**—Because the source is a nonexistent tunnel source, Packet C fails the reverse path forwarding check, and the packet is not forwarded.

## 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.

```

Router R0
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.1/24
set interfaces ge-0/0/2 vlan-tagging
set interfaces ge-0/0/2 unit 0 vlan-id 1
set interfaces ge-0/0/2 unit 0 family inet address 172.16.0.1/16
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set routing-options router-id 10.1.1.1
set routing-options autonomous-system 100
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels T1 source-address 192.0.2.1
set routing-options dynamic-tunnels T1 gre
set routing-options dynamic-tunnels T1 destination-networks 192.0.2.0/24
set routing-options dynamic-tunnels T2 source-address 198.51.100.1
set routing-options dynamic-tunnels T2 gre
set routing-options dynamic-tunnels T2 destination-networks 198.51.100.0/24
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 10.1.1.1
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 20.1.1.1
set protocols bgp group IBGP neighbor 30.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set routing-instances VPN1 instance-type vrf
set routing-instances VPN1 interface ge-0/0/2.0
set routing-instances VPN1 route-distinguisher 100:100
set routing-instances VPN1 vrf-target target:100:1
set routing-instances VPN1 vrf-table-label
set routing-instances VPN1 routing-options forwarding-table ip-tunnel-rpf-check mode
strict
set routing-instances VPN1 protocols bgp group External type external
set routing-instances VPN1 protocols bgp group External family inet unicast
set routing-instances VPN1 protocols bgp group External peer-as 200
set routing-instances VPN1 protocols bgp group External neighbor 172.16.0.1

Router R1
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 0 vlan-id 2
set interfaces ge-0/0/1 unit 0 family inet address 172.18.0.1/16
set interfaces lo0 unit 0 family inet address 20.1.1.1/32
set routing-options router-id 20.1.1.1
set routing-options autonomous-system 100
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels T1 source-address 192.0.2.2

```

```

set routing-options dynamic-tunnels T1 gre
set routing-options dynamic-tunnels T1 destination-networks 192.0.2.0/24
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 20.1.1.1
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 30.1.1.1
set protocols bgp group IBGP neighbor 10.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set routing-instances VPN2 instance-type vrf
set routing-instances VPN2 interface ge-0/0/1.0
set routing-instances VPN2 route-distinguisher 100:200
set routing-instances VPN2 vrf-target target:200:1
set routing-instances VPN2 vrf-table-label

```

```

R2  set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.2/24
    set interfaces ge-0/0/2 vlan-tagging
    set interfaces ge-0/0/2 unit 0 vlan-id 3
    set interfaces ge-0/0/2 unit 0 family inet address 172.17.0.1/16
    set interfaces lo0 unit 0 family inet address 30.1.1.1/32
    set routing-options router-id 30.1.1.1
    set routing-options autonomous-system 100
    set routing-options dynamic-tunnels gre next-hop-based-tunnel
    set routing-options dynamic-tunnels T2 source-address 198.51.100.2
    set routing-options dynamic-tunnels T2 gre
    set routing-options dynamic-tunnels T2 destination-networks 198.51.100.0/24
    set protocols rsvp interface all
    set protocols rsvp interface fxp0.0 disable
    set protocols bgp group IBGP type internal
    set protocols bgp group IBGP local-address 30.1.1.1
    set protocols bgp group IBGP family inet-vpn unicast
    set protocols bgp group IBGP neighbor 20.1.1.1
    set protocols bgp group IBGP neighbor 10.1.1.1
    set protocols ospf traffic-engineering
    set protocols ospf area 0.0.0.0 interface lo0.0 passive
    set protocols ospf area 0.0.0.0 interface all
    set routing-instances VPN3 instance-type vrf
    set routing-instances VPN3 interface ge-0/0/2.0
    set routing-instances VPN3 route-distinguisher 100:300
    set routing-instances VPN3 vrf-target target:300:1
    set routing-instances VPN3 vrf-table-label

```

**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 Router R0:

1. Configure Router R0's interfaces, including the loopback interface.

```

[edit interfaces]
user@R0# set ge-0/0/0 unit 0 family inet address 192.0.2.1/24

```

```
user@R0# set ge-0/0/1 unit 0 family inet address 198.51.100.1/24
```

```
user@R0# set ge-0/0/2 vlan-tagging
user@R0# set ge-0/0/2 unit 0 vlan-id 1
user@R0# set ge-0/0/2 unit 0 family inet address 172.16.0.1/16
```

```
user@R0# set lo0 unit 0 family inet address 10.1.1.1/32
```

2. Assign the router ID and autonomous system number for Router R0.

```
[edit routing-options]
user@R0# set router-id 10.1.1.1
user@R0# set autonomous-system 100
```

3. Configure IBGP peering between the routers.

```
[edit protocols]
user@R0# set bgp group IBGP type internal
user@R0# set bgp group IBGP local-address 10.1.1.1
user@R0# set bgp group IBGP family inet-vpn unicast
user@R0# set bgp group IBGP neighbor 20.1.1.1
user@R0# set bgp group IBGP neighbor 30.1.1.1
```

4. Configure OSPF on all the interfaces of Router R0, excluding the management interface.

```
[edit protocols]
user@R0# set ospf traffic-engineering
user@R0# set ospf area 0.0.0.0 interface lo0.0 passive
user@R0# set ospf area 0.0.0.0 interface all
```

5. Configure RSVP on all the interfaces of Router R0, excluding the management interface.

```
[edit protocols]
user@R0# set rsvp interface all
user@R0# set rsvp interface fxp0.0 disable
```

6. Enable next-hop-based dynamic GRE tunnel configuration on Router R0.

```
[edit routing-options]
user@R0# set dynamic-tunnels gre next-hop-based-tunnel
```

7. Configure the dynamic GRE tunnel parameters from Router R0 to Router R1.

```
[edit routing-options]
user@R0# set dynamic-tunnels T1 source-address 192.0.2.1
user@R0# set dynamic-tunnels T1 gre
user@R0# set dynamic-tunnels T1 destination-networks 192.0.2.0/24
```

8. Configure the dynamic GRE tunnel parameters from Router R0 to Router R2.

```
[edit routing-options]
```

```

user@R0# set dynamic-tunnels T2 source-address 198.51.100.1
user@R0# set dynamic-tunnels T2 gre
user@R0# set dynamic-tunnels T2 destination-networks 198.51.100.0/24

```

9. Configure a virtual routing and forwarding (VRF) routing instance on Router R0, and assign the interface connecting to Host 1 to the VRF instance.

```

[edit routing-instances]
user@R0# set VPN1 instance-type vrf
user@R0# set VPN1 route-distinguisher 100:100
user@R0# set VPN1 vrf-target target:100:1
user@R0# set VPN1 vrf-table-label

```

```

user@R0# set VPN1 interface ge-0/0/2.0

```

10. Configure an external BGP session with Host 1 for the VRF routing instance.

```

[edit routing-instances]
user@R0# set VPN1 protocols bgp group External type external
user@R0# set VPN1 protocols bgp group External family inet unicast
user@R0# set VPN1 protocols bgp group External peer-as 200
user@R0# set VPN1 protocols bgp group External neighbor 172.16.0.1

```

11. Configure anti-spoofing protection for the VRF routing instance on Router R0. This enables reverse path forwarding check for the next-hop-based dynamic tunnels, T1 and T2, on Router 0.

```

[edit routing-instances]
user@R0# set VPN1 routing-options forwarding-table ip-tunnel-rpf-check mode
strict

```

## Results

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

```

user@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 198.51.100.1/24;
    }
  }
}

```

```
}
ge-0/0/2 {
  vlan-tagging;
  unit 0 {
    vlan-id 1;
    family inet {
      address 172.16.0.1/16;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.1.1.1/32;
    }
  }
}

user@R0# show routing-options
router-id 10.1.1.1;
autonomous-system 100;
dynamic-tunnels {
  gre next-hop-based-tunnel;
  T1 {
    source-address 192.0.2.1;
    gre;
    destination-networks {
      192.0.2.0/24;
    }
  }
  T2 {
    source-address 198.51.100.1;
    gre;
    destination-networks {
      198.51.100.0/24;
    }
  }
}

user@R0# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  group IBGP {
    type internal;
    local-address 10.1.1.1;
    family inet-vpn {
      unicast;
    }
    neighbor 20.1.1.1;
    neighbor 30.1.1.1;
  }
}
```

```

ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface all;
  }
}

user@R0# show routing-instances
VPN1 {
  instance-type vrf;
  interface ge-0/0/2.0;
  route-distinguisher 100:100;
  vrf-target target:100:1;
  vrf-table-label;
  routing-options {
    forwarding-table {
      ip-tunnel-rpf-check {
        mode strict;
      }
    }
  }
  protocols {
    bgp {
      group External {
        type external;
        family inet {
          unicast;
        }
        peer-as 200;
        neighbor 172.16.0.1;
      }
    }
  }
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying Basic Configuration on page 763](#)
- [Verifying Dynamic Tunnel Configuration on page 764](#)
- [Verifying Anti-Spoofing Protection Configuration on page 765](#)

### Verifying Basic Configuration

**Purpose** Verify the OSPF and BGP peering status between the Router R0 and Routers R1 and R2.

**Action** From operational mode, run the **show ospf neighbor** and **show bgp summary** commands.

```

user@R0> show ospf neighbor

```

Address	Interface	State	ID	Pri	Dead
192.0.2.2	ge-0/0/0.0	Full	20.1.1.1	128	32
198.51.100.2	ge-0/0/1.0	Full	30.1.1.1	128	32

```

user@R0> show bgp summary
Groups: 2 Peers: 3 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0
0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
20.1.1.1 100 182 178 0 0 1:20:27
Establ
bgp.13vpn.0: 0/0/0/0
30.1.1.1 100 230 225 0 0 1:41:51
Establ
bgp.13vpn.0: 0/0/0/0
172.16.0.1 200 0 0 0 0 1:42:08
Establ

```

**Meaning** The OSPF and BGP sessions are up and running between the Routers R0, R1, and R2.

### Verifying Dynamic Tunnel Configuration

**Purpose** Verify the status of the next-hop-based dynamic GRE tunnels between the Router R0 and Routers R1 and R2.



**Action** From operational mode, run the **show route table inet.3**, and the **show dynamic-tunnels database terse** commands.

```
user@R0> show route table inet.3

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

192.0.2.0/24      *[Tunnel/300] 01:47:57
                  Tunnel
192.0.2.2/24      *[Tunnel/300] 01:47:57
                  Tunnel Composite

198.51.100.0/24   *[Tunnel/300] 01:47:57
                  Tunnel
198.51.100.2/24   *[Tunnel/300] 01:47:57
                  Tunnel Composite

user@R0> show dynamic-tunnels database terse
Table: inet.3

Destination-network: 192.0.2.0/24
Destination      Source      Next-hop      Type      Status
192.0.2.2/24     192.0.2.1   0xb395e70     nhid 612   gre      Up

Destination-network: 198.51.100.0/24
Destination      Source      Next-hop      Type      Status
198.51.100.2     198.51.100.1 0xb395e70     nhid 612   gre      Up
```

**Meaning** The two next-hop-based dynamic GRE tunnels, Tunnel 1 and Tunnel 2, are up.

### Verifying Anti-Spoofing Protection Configuration

**Purpose** Verify that the reverse path forwarding check has been enabled on the VRF routing instance on Router R0.

**Action** From the operational mode, run the **show krt table VPN1.inet.0 detail**.

```
user@R0> show krt table VPN1.inet.0 detail
KRT tables:
VPN1.inet.0          : GF: 1 krt-index: 8      ID: 0 kernel-id: 8
    flags: (null)
    tunnel rpf config data : enable, strict, filter [0], 0x2
    tunnel rpf tlv data : enable, strict, filter [0], 0x4
    unicast reverse path: disabled
    fast-reroute-priority: 0
    Permanent NextHops
      Multicast      : 0 Broadcast : 0
      Receive        : 0 Discard   : 0
      Multicast Discard: 0 Reject   : 0
      Local          : 0 Deny      : 0
      Table          : 0
```

**Meaning** The configured reverse path forwarding check is enabled on the VRF routing instance in the strict mode.

**Related Documentation**

- [Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview on page 753](#)
- [ip-tunnel-rpf-check on page 1080](#)

## PART 7

# Configuring Load Balancing in Layer 3 VPNs

- [Configuring Load Balancing in Layer 3 VPNs on page 769](#)
- [Enabling BGP PIC for Layer 3 VPNs on page 791](#)



# Configuring Load Balancing in Layer 3 VPNs

- [VPN Per-Packet Load Balancing on page 769](#)
- [Load Balancing and IP Header Filtering for Layer 3 VPNs on page 770](#)
- [Layer 3 VPN Load Balancing Using IP Header Filtering on page 771](#)
- [Configuring Protocol-Independent Load Balancing in Layer 3 VPNs on page 787](#)

## VPN Per-Packet Load Balancing

---

By default, when there are multiple equal-cost paths to the same destination for the active route, the Junos OS software uses a hash algorithm to select one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes, this selection process (using the same hash algorithm) is repeated to choose the best single next-hop address using the same hash algorithm.

Alternatively, you can configure the Junos OS software to spread the VPN traffic across the multiple valid paths between PE devices. This feature is called per-packet load balancing. VPN traffic load balancing is only possible when more than one valid path is available. You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. In addition to increasing the volume of traffic you can send between VPN devices, you can configure per-packet load balancing to optimize traffic flows across multiple paths.

Traffic is distributed across multiple valid paths by running a hash algorithm on various elements of the route, such as the MPLS label or the destination address. The following tables describe how the load balancing hash algorithm is run on routes at the ingress router and at the transit and egress routers. The route elements used by the hash algorithm vary depending on VPN application. If Junos OS encounters an S-bit set to 1 (indicating the bottom of the stack), it does not apply the hash algorithm any further.

**Table 12: Ingress Router Hashing**

Application	Ingress Logical Interface	MPLS Labels	Source and Destination MAC Addresses	Reordering and Flow Separation Risk	Disable Control Word	IP (Source/Destination Address and Port, Protocol)
Layer 2 VPNs and Layer 2 Circuits configured with CCC	Yes	Yes	No	Yes (if the data is variable, for example ATM)	Yes	N/A
Layer 2 VPNs and Layer 2 Circuits configured with TCC	Yes	Yes	No	Yes (if the data is variable, for example ATM)	Yes	N/A
Layer 3 VPNs and IPv4 or IPv6 RIBs	Yes	No	No	No	No	Yes
VPLS	Yes	No	Yes	No	No	Yes

**Table 13: Transit and Egress Router Hashing**

Application	Ingress Logical Interface	MPLS Labels (up to 3 and the S-bit is set to 1)	Reordering and Flow Separation Risk	IP (Source/Destination Address and Port, Protocol)
Layer 2 VPNs and Layer 2 Circuits configured with CCC	Yes	Yes	No	No
Layer 2 VPNs and Layer 2 Circuits configured with TCC	Yes	Yes	No	Yes
Layer 3 VPNs and IPv4 or IPv6 RIBs	Yes	Yes	No	Yes
VPLS	Yes	Yes for known unicast traffic  No for broadcast, unicast unknown, and multicast traffic	No	No

**Related Documentation** • [Configuring Protocol-Independent Load Balancing in Layer 3 VPNs on page 787](#)

## Load Balancing and IP Header Filtering for Layer 3 VPNs

You can now simultaneously enable both load balancing of traffic across both internal and external BGP paths and filtering of traffic based on the IP header. This enables you to configure filters and policers at the egress PE router for traffic that is simultaneously

being load-balanced across both internal and external BGP paths. This feature is available only on the M120 router, M320 router, MX Series routers, and T Series routers.

To enable these features on a Layer 3 VPN routing instance, include the **vpn-unequal-cost equal-external-internal** statement at the **[edit routing-instances routing-instance-name routing-options multipath]** hierarchy level and the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy level.

If you issue the **show route detail** command, you can discover whether or not a route is being load-balanced (equal-external-internal) and what its interface index is.

If you have also configured fast reroute, please be aware of the following behavior:

- If an IBGP path goes down, it could be replaced by either an active EBGP path or an active IBGP path.
- If an EBGP path goes down, it can only be replaced by another active EBGP path. This prevents the forwarding of core-facing interface traffic to an IBGP destination.



**NOTE:** You can include the **vpn-unequal-cost equal-external-internal** statement and the **l3vpn** statement at the **[edit routing-options forwarding-options chained-composite-next-hop ingress]** hierarchy level simultaneously. However, if you do this, EBGP does not work. This means that when there are both paths with chained nexthops and paths with nonchained nexthops as candidates for EBGP equal-cost multipath (ECMP), the paths using chained nexthops are excluded. In a typical case, the excluded paths are the internal paths.

## Layer 3 VPN Load Balancing Using IP Header Filtering

- [Layer 3 VPN Load Balancing Overview on page 771](#)
- [Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering on page 772](#)

### Layer 3 VPN Load Balancing Overview

The load balancing feature allows a device to divide incoming and outgoing traffic along multiple paths in order to reduce congestion in the network. Load balancing improves the utilization of various network paths, and provides more effective network bandwidth.

When multiple protocols are in use, the device uses the route preference value (also known as the administrative distance value) to select a route. While using a single routing protocol, the router chooses the path with the lowest cost (or metric) to the destination. If the device receives and installs multiple paths with the same route preference and same cost to a destination, load balancing must be configured.

In a network with both internal and external BGP paths installed among devices in different autonomous systems, BGP selects only a single best path by default, and does not perform load balancing. A Layer 3 VPN with internal and external BGP paths uses the

**multipath** statement for protocol-independent load balancing. When you include the **multipath** statement in a routing instance, protocol-independent load balancing is applied to the default routing table for that routing instance. By using the **vpn-unequal-cost** statement, protocol-independent load balancing is applied to VPN routes. By using the **equal-external-internal** statement, protocol-independent load balancing is applied to both internal and external BGP paths and can be configured in conjunction with IP header filtering (enabled with the **vrf-table-label** statement).

- See Also**
- [Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering on page 772](#)

## Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering

This example shows how to configure load balancing in a Layer 3 VPN (with internal and external BGP paths) while simultaneously using IP header filtering.

- [Requirements on page 772](#)
- [Overview on page 772](#)
- [Configuration on page 774](#)
- [Verification on page 782](#)

### Requirements

---

This example requires the following hardware and software components:

- M Series Multiservice Edge Routers (M120 and M320 only), MX Series 3D Universal Edge Routers, T Series Core Routers, or PTX Series Transport Routers.
- Junos OS Release 12.1 or later

### Overview

---

The following example shows how to configure load balancing while simultaneously using IP header filtering in a Layer 3 VPN.



**NOTE:** This example demonstrates how load balancing and IP header filtering work together. The testing of IP header filtering is out of the scope of this example.

The Junos OS BGP provides a multipath feature that allows load balancing between peers in the same or different autonomous systems (ASs). This example uses the **equal-external-internal** statement at the **[edit routing-instances instance-name routing-options multipath vpn-unequal-cost]** hierarchy level to perform load balancing. The **vrf-table-label** statement is configured at the **[edit routing-instances instance-name]** hierarchy level to enable IP header filtering.

```
[edit]
routing-instances {
```



```

instance-name {
  vrf-table-label;
  routing-options {
    multipath {
      vpn-unequal-cost {
        equal-external-internal;
      }
    }
  }
}

```



**NOTE:** These statements are available only in the context of a routing instance.

In this example, Device CE1 is in AS1 and connected to Device PE1. Devices PE1, PE2, PE3, and P are in AS2. Device CE2 is connected to Devices PE2 and PE3 and is in AS3. Device CE3 is connected to Device PE3 and is in AS4. BGP and MPLS are configured through the network. OSPF is the interior gateway protocol (IGP) that is used in this network.

The configuration for Devices PE1, PE2, and PE3 includes the `equal-external-internal` statement at the `[edit routing-instances instance-name routing-options multipath vpn-unequal-cost]` hierarchy level to enable load balancing in the network. IP header filtering is enabled when the `vrf-table-label` statement is configured at the `[edit routing-instances instance-name]` hierarchy level on the PE devices.

Figure 75 on page 773 shows the topology used in this example.

*Figure 75: Layer 3 VPN Load Balancing Using IP Header Filtering*

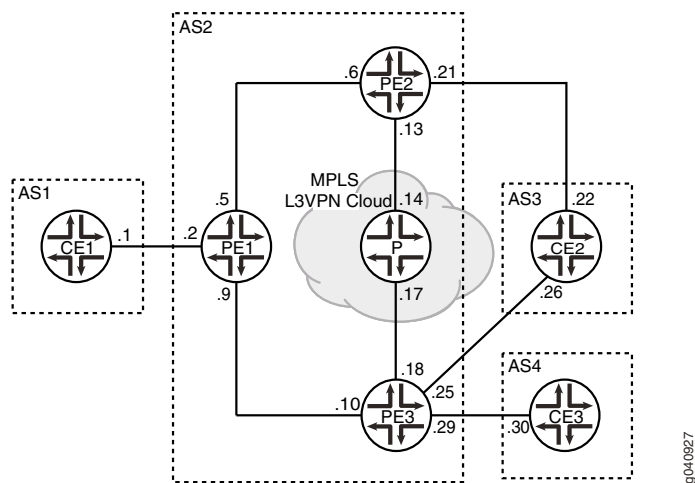


Table 14 on page 774 shows the list of IP addresses used in this example for quick reference.

Table 14: Device IP Address Quick Reference

Device	AS	Device ID	Device Interface Units	Device Interface Unit IPs
CE1	1	192.0.2.1/24	Unit 1	10.1.1.1/30
PE1	2	192.0.2.2/24	Unit 2	10.1.1.2/30
			Unit 5	10.1.2.5/30
			Unit 9	10.1.3.9/30
PE2	2	192.0.2.3/24	Unit 6	10.1.2.6/30
			Unit 13	10.1.4.13/30
			Unit 21	10.1.6.21/30
PE3	2	192.0.2.4/24	Unit 10	10.1.3.10/30
			Unit 18	10.1.5.18/30
			Unit 25	10.1.7.25/30
			Unit 29	10.1.8.29/30
P	2	192.0.2.5/24	Unit 14	10.1.4.14/30
			Unit 17	10.1.5.17/30
CE2	3	192.0.2.6/24	Unit 22	10.1.6.22/30
			Unit 26	10.1.7.26/30
CE3	4	192.0.2.7/24	Unit 30	10.1.8.30/30



**NOTE:** This example was tested using logical systems (logical routers). Therefore all the physical interfaces in the example are the same and the configuration is done on separate logical interfaces. In an non-test network, you will use separate physical routers and separate physical interfaces for the connections to other devices.

### 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 CE1

```
set interfaces ge-2/1/10 unit 1 family inet address 10.1.1.1/30
set interfaces ge-2/1/10 unit 1 family mpls
set interfaces ge-2/1/10 unit 1 description toPE1
set interfaces lo0 unit 4 family inet address 192.0.2.1/24
set routing-options router-id 192.0.2.1
set routing-options autonomous-system 1
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 export send-direct
set protocols bgp group toPE1 peer-as 2
set protocols bgp group toPE1 neighbor 10.1.1.2
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

#### Device PE1

```
set interfaces ge-2/1/10 unit 2 family inet address 10.1.1.2/30
set interfaces ge-2/1/10 unit 2 family mpls
set interfaces ge-2/1/10 unit 2 description toCE1
set interfaces ge-2/1/10 unit 5 family inet address 10.1.2.5/30
set interfaces ge-2/1/10 unit 5 family mpls
set interfaces ge-2/1/10 unit 5 description toPE2
set interfaces ge-2/1/10 unit 9 family inet address 10.1.3.9/30
set interfaces ge-2/1/10 unit 9 family mpls
set interfaces ge-2/1/10 unit 9 description toPE3
set interfaces lo0 unit 5 family inet address 192.0.2.2/24
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.5 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/1/10.9 metric 10
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal local-address 192.0.2.2
set protocols bgp group toInternal neighbor 192.0.2.3
set protocols bgp group toInternal neighbor 192.0.2.4
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 2
set routing-options forwarding-table export lb
set routing-instances purple instance-type vrf
set routing-instances purple interface ge-2/1/10.2
set routing-instances purple route-distinguisher 2:1
set routing-instances purple vrf-target target:2:1
set routing-instances purple vrf-table-label
set routing-instances purple protocols bgp group toCE1 type external
set routing-instances purple protocols bgp group toCE1 peer-as 1
set routing-instances purple protocols bgp group toCE1 neighbor 10.1.1.1
set routing-instances purple routing-options multipath vpn-unequal-cost
    equal-external-internal
set policy-options policy-statement lb then load-balance per-packet
```

#### Device PE2

```
set interfaces ge-2/1/10 unit 6 family inet address 10.1.2.6/30
set interfaces ge-2/1/10 unit 6 family mpls
set interfaces ge-2/1/10 unit 6 description toPE1
set interfaces ge-2/1/10 unit 13 family inet address 10.1.4.13/30
set interfaces ge-2/1/10 unit 13 family mpls
set interfaces ge-2/1/10 unit 13 description toP
set interfaces ge-2/1/10 unit 21 family inet address 10.1.6.21/30
set interfaces ge-2/1/10 unit 21 family mpls
set interfaces ge-2/1/10 unit 21 description toCE2
set interfaces lo0 unit 6 family inet address 192.0.2.3/24
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.6 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/1/10.13 metric 5
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal local-address 192.0.2.3
set protocols bgp group toInternal neighbor 192.0.2.2
set protocols bgp group toInternal neighbor 192.0.2.4
set routing-options router-id 192.0.2.3
set routing-options autonomous-system 2
set routing-options forwarding-table export lb
set routing-instances purple instance-type vrf
set routing-instances purple interface ge-2/1/10.21
set routing-instances purple route-distinguisher 2:1
set routing-instances purple vrf-target target:2:1
set routing-instances purple vrf-table-label
set routing-instances purple protocols bgp group toCE2 type external
set routing-instances purple protocols bgp group toCE2 peer-as 3
set routing-instances purple protocols bgp group toCE2 neighbor 10.1.6.22
set routing-instances purple routing-options multipath vpn-unequal-cost
    equal-external-internal
set policy-options policy-statement lb then load-balance per-packet
```

#### Device PE3

```
set interfaces ge-2/1/10 unit 10 family inet address 10.1.3.10/30
set interfaces ge-2/1/10 unit 10 family mpls
set interfaces ge-2/1/10 unit 10 description toPE1
set interfaces ge-2/1/10 unit 18 family inet address 10.1.5.18/30
set interfaces ge-2/1/10 unit 18 family mpls
set interfaces ge-2/1/10 unit 18 description toP
set interfaces ge-2/1/10 unit 25 family inet address 10.1.7.25/30
set interfaces ge-2/1/10 unit 25 family mpls
set interfaces ge-2/1/10 unit 25 description toCE2
set interfaces ge-2/1/10 unit 29 family inet address 10.1.8.29/30
set interfaces ge-2/1/10 unit 29 family mpls
set interfaces ge-2/1/10 unit 29 description toCE3
set interfaces lo0 unit 7 family inet address 192.0.2.4/24
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.10 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/1/10.18 metric 5
set protocols bgp group toInternal type internal
```

```

set protocols bgp group toInternal local-address 192.0.2.4
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal family route-target
set protocols bgp group toInternal neighbor 192.0.2.2
set protocols bgp group toInternal neighbor 192.0.2.3
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 2
set routing-options forwarding-table export lb
set routing-instances purple instance-type vrf
set routing-instances purple interface ge-2/1/10.25
set routing-instances purple interface ge-2/1/10.29
set routing-instances purple route-distinguisher 2:1
set routing-instances purple vrf-target target:2:1
set routing-instances purple vrf-table-label
set routing-instances purple protocols bgp group toCE2 type external
set routing-instances purple protocols bgp group toCE2 peer-as 3
set routing-instances purple protocols bgp group toCE2 neighbor 10.1.7.26
set routing-instances purple protocols bgp group toCE3 type external
set routing-instances purple protocols bgp group toCE3 peer-as 4
set routing-instances purple protocols bgp group toCE3 neighbor 10.1.8.30
set routing-instances purple routing-options multipath vpn-unequal-cost
    equal-external-internal
set policy-options policy-statement lb then load-balance per-packet

```

#### Device P

```

set interfaces ge-2/1/10 unit 14 family inet address 10.1.4.14/30
set interfaces ge-2/1/10 unit 14 family mpls
set interfaces ge-2/1/10 unit 14 description toPE2
set interfaces ge-2/1/10 unit 17 family inet address 10.1.5.17/30
set interfaces ge-2/1/10 unit 17 family mpls
set interfaces ge-2/1/10 unit 17 description toPE3
set interfaces lo0 unit 8 family inet address 192.0.2.5/24
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.8 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.14 metric 5
set protocols ospf area 0.0.0.0 interface ge-2/1/10.17 metric 5
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 2

```

#### Device CE2

```

set interfaces ge-2/1/10 unit 22 family inet address 10.1.6.22/30
set interfaces ge-2/1/10 unit 22 family mpls
set interfaces ge-2/1/10 unit 22 description toPE2
set interfaces ge-2/1/10 unit 26 family inet address 10.1.7.26/30
set interfaces ge-2/1/10 unit 26 family mpls
set interfaces ge-2/1/10 unit 26 description toPE3
set interfaces lo0 unit 6 family inet address 192.0.2.6/24
set routing-options router-id 192.0.2.6
set routing-options autonomous-system 3
set protocols bgp group toAS2 type internal
set protocols bgp group toAS2 export send-direct
set protocols bgp group toAS2 peer-as 2
set protocols bgp group toAS2 neighbor 10.1.6.21
set protocols bgp group toAS2 neighbor 10.1.7.25

```

```
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

#### Device CE3

```
set interfaces ge-2/1/10 unit 30 family inet address 10.1.8.30/30
set interfaces ge-2/1/10 unit 30 family mpls
set interfaces ge-2/1/10 unit 30 description toPE3
set interfaces lo0 unit 7 family inet address 192.0.2.7/24
set routing-options router-id 192.0.2.7
set routing-options autonomous-system 4
set protocols bgp group toPE3 type internal
set protocols bgp group toPE3 export send-direct
set protocols bgp group toPE3 peer-as 2
set protocols bgp group toPE3 neighbor 10.1.8.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

**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 unequal-cost load balancing across the VPN setup:

1. Configure the router ID on Device CE1, and assign the device to its autonomous system.

```
[edit routing-options]
user@CE1# set router-id 192.0.2.1
user@CE1# set autonomous-system 1
```

Similarly, configure all other devices.

2. Configure BGP groups for traffic through the entire network.
  - a. Configure the BGP group for traffic to and from the MPLS network (CE devices).

```
[edit protocols bgp group toPE1]
user@CE1# set type external
user@CE1# set peer-as 2
user@CE1# set neighbor 10.1.1.2
```

- b. Configure similar BGP groups (**toAS2** and **toPE3**) on Devices CE2 and CE3 by modifying the **peer-as** and **neighbor** statements accordingly.
- c. Configure the BGP group for traffic through the MPLS network (PE devices).

```
[edit protocols bgp group toInternal]
user@PE1# set type internal
user@PE1# set family inet-vpn unicast
user@PE1# set local-address 192.0.2.2
user@PE1# set neighbor 192.0.2.3
user@PE1# set neighbor 192.0.2.4
```

- d. Configure the same BGP group (**toInternal**) on Devices PE2 and PE3 by modifying the **local-address** and **neighbor** statements accordingly.

3. Configure a routing policy for exporting routes to and from the MPLS network (**send-direct** policy) and a policy for load balancing traffic network across the MPLS network (**lb** policy).

- a. Configure a policy (**send-direct**) for exporting routes from the routing table into BGP on Device CE1.

```
[edit policy-options policy-statement send-direct]
user@CE1# set from protocol direct
user@CE1# set then accept

[edit protocols bgp group toPE1]
user@CE1# set export send-direct
```

Similarly, configure the **send-direct** policy on Devices CE2 and CE3.

- b. Configure a policy (**lb**) for exporting routes from the routing table into the forwarding table on Device PE1.

The **lb** policy configures per-packet load balancing, which ensures that all next-hop addresses for a destination are installed in the forwarding table.

```
[edit policy-options policy-statement lb]
user@PE1# set then load-balance per-packet

[edit routing-options]
user@PE1# set forwarding-table export lb
```

Similarly, configure the **lb** policy on Devices PE2, and PE3.

4. Configure the following:
  - a. Configure the routing instance on the PE devices for exporting routes through the autonomous systems.
  - b. Include the **equal-external-internal** statement at the **[edit routing-instances instance-name routing-options multipath vpn-unequal-cost]** hierarchy level to enable load balancing in the network.
  - c. Include the **vrf-table-label** statement at the **[edit routing-instances instance-name]** hierarchy level for filtering traffic prior to exiting the egress device (Device CE3).

#### Device PE1

```
[edit routing-instances purple]
user@PE1# set instance-type vrf
user@PE1# set interface ge-2/1/10.2
user@PE1# set route-distinguisher 2:1
user@PE1# set vrf-target target:2:1
user@PE1# set vrf-table-label
user@PE1# set protocols bgp group toCE1 type external
user@PE1# set protocols bgp group toCE1 peer-as 1
user@PE1# set protocols bgp group toCE1 neighbor 10.1.1.1
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```

#### Device PE2

```
[edit routing-instances purple]
user@PE2# set instance-type vrf
```

```

user@PE2# set interface ge-2/1/10.21
user@PE2# set route-distinguisher 2:1
user@PE2# set vrf-target target:2:1
user@PE2# set vrf-table-label
user@PE2# set protocols bgp group toCE2 type external
user@PE2# set protocols bgp group toCE2 peer-as 3
user@PE2# set protocols bgp group toCE2 neighbor 10.1.6.22
user@PE2# set routing-options multipath vpn-unequal-cost equal-external-internal

```

#### Device PE3

```

[edit routing-instances purple]
user@PE3# set instance-type vrf
user@PE3# set interface ge-2/1/10.25
user@PE3# set interface ge-2/1/10.29
user@PE3# set route-distinguisher 2:1
user@PE3# set vrf-target target:2:1
user@PE3# set vrf-table-label
user@PE3# set protocols bgp group toCE2 type external
user@PE3# set protocols bgp group toCE2 peer-as 3
user@PE3# set protocols bgp group toCE2 neighbor 10.1.7.26
user@PE3# set protocols bgp group toCE3 type external
user@PE3# set protocols bgp group toCE3 peer-as 4
user@PE3# set protocols bgp group toCE3 neighbor 10.1.8.30
user@PE3# set routing-options multipath vpn-unequal-cost equal-external-internal

```

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

```

user@PE3# show interfaces
ge-2/1/10 {
  unit 10 {
    description toPE1;
    family inet {
      address 10.1.3.10/30;
    }
    family mpls
  }
  unit 18 {
    description toP;
    family inet {
      address 10.1.5.18/30;
    }
    family mpls
  }
  unit 25 {
    description toCE2;
    family inet {
      address 10.1.7.25/30;
    }
    family mpls
  }
}

```



```
unit 29 {
  description toCE3;
  family inet {
    address 10.1.8.29/30;
  }
  family mpls
}
}
lo0 {
  unit 7 {
    family inet {
      address 192.0.2.4/24;
    }
  }
}

user@PE3# show protocols
mpls {
  interface all;
}
bgp {
  group toInternal {
    type internal;
    local-address 192.0.2.4;
    family inet {
      unicast;
    }
    family inet-vpn {
      unicast;
    }
    family route-target;
    neighbor 192.0.2.2;
    neighbor 192.0.2.3;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.7 {
      passive;
    }
    interface ge-2/1/10.10 {
      metric 10;
    }
    interface ge-2/1/10.18 {
      metric 5;
    }
  }
}
ldp {
  interface all;
}

user@PE3# show policy-options
policy-statement lb {
  then {
    load-balance per-packet;
  }
}
```

```
}  
user@PE3# show routing-instances  
purple {  
  instance-type vrf;  
  interface ge-2/1/10.25;  
  interface ge-2/1/10.29;  
  route-distinguisher 2:1;  
  vrf-target target:2:1;  
  vrf-table-label;  
  routing-options {  
    multipath {  
      vpn-unequal-cost equal-external-internal;  
    }  
  }  
}  
protocols {  
  bgp {  
    group toCE2 {  
      type external;  
      peer-as 3;  
      neighbor 10.1.7.26;  
    }  
    group toCE3 {  
      type external;  
      peer-as 4;  
      neighbor 10.1.8.30;  
    }  
  }  
}  
}  
user@PE3# show routing-options  
router-id 192.0.2.4;  
autonomous-system 2;  
forwarding-table {  
  export lb;  
}
```

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

---

## Verification

Confirm that the configuration is working properly.

- [Verifying BGP on page 782](#)
- [Verifying Load Balancing on page 784](#)
- [Verifying Load Balancing While Using IP Header Filtering on page 786](#)

### *Verifying BGP*

**Purpose** Verify that BGP is working.

**Action** From operational mode, run the **show route protocol bgp** command.

```

user@PE3> show route protocol bgp

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

purple.inet.0: 9 destinations, 14 routes (9 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

192.0.2.1/24      *[BGP/170] 04:47:14, localpref 100, from 192.0.2.2
                  AS path: 1 I
                  > to 10.1.3.9 via ge-2/1/10.10, Push 16
192.0.2.6/24      @[BGP/170] 00:13:28, localpref 100
                  AS path: 3 I
                  > to 10.1.7.26 via ge-2/1/10.25
                  [BGP/170] 00:10:36, localpref 100, from 192.0.2.3
                  AS path: 3 I
                  > to 10.1.5.17 via ge-2/1/10.18, Push 16, Push 299776(top)
192.0.2.7/24      *[BGP/170] 00:10:56, localpref 100
                  AS path: 4 I
                  > to 10.1.8.30 via ge-2/1/10.29
10.1.1.0/30       *[BGP/170] 04:47:14, localpref 100, from 192.0.2.2
                  AS path: I
                  > to 10.1.3.9 via ge-2/1/10.10, Push 16
10.1.6.20/30      *[BGP/170] 04:47:03, localpref 100, from 192.0.2.3
                  AS path: I
                  > to 10.1.5.17 via ge-2/1/10.18, Push 16, Push 299776(top)
                  [BGP/170] 00:13:28, localpref 100
                  AS path: 3 I
                  > to 10.1.7.26 via ge-2/1/10.25
10.1.7.24/30      [BGP/170] 00:13:28, localpref 100
                  AS path: 3 I
                  > to 10.1.7.26 via ge-2/1/10.25
10.1.8.28/30      [BGP/170] 00:10:56, localpref 100
                  AS path: 4 I
                  > to 10.1.8.30 via ge-2/1/10.29

mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:1:192.0.2.1/24  *[BGP/170] 04:47:14, localpref 100, from 192.0.2.2
                  AS path: 1 I
                  > to 10.1.3.9 via ge-2/1/10.10, Push 16
2:1:192.0.2.6/24  *[BGP/170] 00:10:36, localpref 100, from 192.0.2.3
                  AS path: 3 I
                  > to 10.1.5.17 via ge-2/1/10.18, Push 16, Push 299776(top)
2:1:10.1.1.0/30   *[BGP/170] 04:47:14, localpref 100, from 192.0.2.2
                  AS path: I
                  > to 10.1.3.9 via ge-2/1/10.10, Push 16
2:1:10.1.6.20/30  *[BGP/170] 04:47:03, localpref 100, from 192.0.2.3

```

The output lists the BGP routes installed into the routing table. The lines of output that start with **192.0.2.1/24**, **10.1.1.0/30**, and **2:1:192.0.2.1/24** show the BGP routes to Device CE1, which is in AS1. The lines of output that start with **192.0.2.6/24**, **2:1:192.0.2.6/24**, and **2:1:10.1.6.20/30** show the BGP routes to Device CE2, which is in AS3. The line of output that starts with **192.0.2.7/24** shows the BGP route to Device CE3, which is in AS4.

**Meaning** BGP is functional in the network.

### Verifying Load Balancing

**Purpose** Verify that forwarding is taking place in both directions by checking:

- If both next hops are installed in the forwarding table for a route.
- If external BGP routes are installed in the forwarding table for a route.

**Action** From operational mode, run the **show route forwarding-table** and **show route forwarding-table destination <destination IP>** commands.

user@PE3> show route forwarding-table

```
Router: PE3
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
192.0.2.2/24     user  1 10.1.3.9      ucst  999   8 ge-2/1/10.10
192.0.2.3/24     user  1 10.1.5.17     ucst 1243  12 ge-2/1/10.18
192.0.2.4/24     intf  0 192.0.2.4     locl  895   1
192.0.2.5/24     user  1 10.1.5.17     ucst 1243  12 ge-2/1/10.18
10.1.2.4/30      user  0
                  10.1.3.9      ucst  999   8 ge-2/1/10.10
                  10.1.5.17     ucst 1243  12 ge-2/1/10.18
10.1.3.8/30      intf  0
10.1.3.8/32      dest  0 10.1.3.8      recv  897   1 ge-2/1/10.10
10.1.3.9/32      dest  0 0.6.80.3.0.21.59.d.c5.d9.0.21.59.d.c5.da.8.0
                  ucst  999   8 ge-2/1/10.10
10.1.3.10/32     intf  0 10.1.3.10     locl  898   2
10.1.3.10/32     dest  0 10.1.3.10     locl  898   2
10.1.3.11/32     dest  0 10.1.3.11     bcst  896   1 ge-2/1/10.10
10.1.4.12/30     user  0 10.1.5.17     ucst 1243  12 ge-2/1/10.18
10.1.5.16/30     intf  0
10.1.5.16/32     dest  0 10.1.5.16     recv  901   1 ge-2/1/10.18
10.1.5.17/32     dest  0 0.e.80.3.0.21.59.d.c5.d9.0.21.59.d.c5.da.8.0
                  ucst 1243  12 ge-2/1/10.18
10.1.5.18/32     intf  0 10.1.5.18     locl  902   2
10.1.5.18/32     dest  0 10.1.5.18     locl  902   2
10.1.5.19/32     dest  0 10.1.5.19     bcst  900   1 ge-2/1/10.18
203.0.113.0/24   perm  2
203.0.113.1/24   perm  0 203.0.0.1     mcst  576   3
203.0.113.5/24   user  1 203.0.0.5     mcst  576   3
255.255.255.255/32 perm  0
                  bcst  577   1
```

Router: PE3

Routing table: \_\_master.anon\_\_.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	909	1	
0.0.0.0/32	perm	0		dscd	907	1	
203.0.113.0/24	perm	0		mdsc	908	1	
203.0.113.1/24	perm	0	203.0.0.1	mcst	904	1	
255.255.255.255/32	perm	0		bcst	905	1	

Router: PE3

Routing table: purple.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	918	1	
0.0.0.0/32	perm	0		dscd	916	1	
192.0.2.1/24	user	0		indr	1048576	3	
			10.1.3.9	Push 16	1187	2	ge-2/1/10.10
192.0.2.6/24	user	0		ulst	1048587	2	
			10.1.7.26	ucst	1239	4	ge-2/1/10.25
			10.1.5.17	indr	1048579	3	
				Push 16, Push	299776(top)	1306	
2 ge-2/1/10.18							
192.0.2.7/24	user	0	10.1.8.30	ucst	1299	4	ge-2/1/10.29
299808(S=0)	user	0	10.1.5.17	Pop	1304	2	ge-2/1/10.18
...							

In the **default.inet** routing table, which is the forwarding table, the line of output that starts with **10.1.2.4/30** shows that for a route to Device PE2 in the same AS, two next hops are installed in the table: **10.1.3.9** and **10.1.5.17**.

In the **purple.inet** routing table, which is the external routing table, the line of output that starts with **192.0.2.6/24** shows that for a route to Device CE2 in AS3, an internal next hop of **10.1.5.17** and an external next hop of **10.1.7.26** are installed in the table. This indicates that both internal and external BGP routes are operational in the network.

user@PE3> show route forwarding-table destination 10.1.2.6

Router: PE3

Routing table: default.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
10.1.2.4/30	user	0		ulst	1048580	2	
			10.1.3.9	ucst	999	8	ge-2/1/10.10
			10.1.5.17	ucst	1243	12	ge-2/1/10.18

Router: PE3

Routing table: \_\_master.anon\_\_.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	909	1	

Router: PE3

Routing table: purple.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	918	1	

The line of output that starts with **10.1.2.4/30** shows that for a route from Device PE3 to Device PE2 in the same AS, two next hops are installed in the table: **10.1.3.9** through the **ge-2/1/10.10** interface, and **10.1.5.17** through the **ge-2/1/10.18** interface.

**Meaning** Multiple next hops for a route, including external BGP routes, are installed in the forwarding tables.

#### *Verifying Load Balancing While Using IP Header Filtering*

**Purpose** Verify that filtered traffic reaches the egress CE devices after load balancing has been configured on the PE devices.

**Action** Configure a firewall filter on Device PE3 on the interface connecting to Device CE2.

```
[edit firewall family inet filter filterPE3 term a]
user@PE3# set from protocol tcp
user@PE3# set from source-port-except bgp
user@PE3# set from destination-port-except bgp
user@PE3# set then count filterPE3
user@PE3# set then accept

[edit firewall family inet filter filterPE3 term b]
user@PE3# set then accept

[edit interfaces ge-2/1/10 unit 25]
user@PE3# set family inet filter output filterPE3
```

Similarly, configure a firewall filter on Device PE3 on the interface facing Device CE3, and another on Device PE2 on the interface facing Device CE2.

Count the packets exiting the egress interfaces on Devices PE2 and PE3 by using the **show firewall filter <filter name> counter <counter name>** operational mode command. The output confirms if load balancing takes place with IP header filtering configured (enabled by the **vrf-table-label** statement). If all transmitted packets have been load-balanced between the paths PE3->CE2, PE3->CE3, and PE2->CE2, then it means that the IP header filtering feature works in a load-balanced Layer 3 network.

You can clear the counter by using the **clear firewall filter <filter name> counter <counter name>** operational mode command.

**Meaning** Load balancing takes place with IP header filtering configured.

**See Also**

- [Configuring Load Balancing for Layer 3 VPNs on page 787](#)
- [Example: Load Balancing BGP Traffic](#)
- [Load Balancing and IP Header Filtering for Layer 3 VPNs on page 770](#)

## Configuring Protocol-Independent Load Balancing in Layer 3 VPNs

Protocol-independent load balancing for Layer 3 VPNs allows the forwarding next hops of both the active route and alternative paths to be used for load balancing. Protocol-independent load balancing works in conjunction with Layer 3 VPNs. It supports the load balancing of VPN routes independently of the assigned route distinguisher. When protocol-independent load balancing is enabled, both routes to other PE routers and routes to directly connected CE routers are load-balanced.

When load-balancing information is created for a given route, the active path is marked as **Routing Use Only** in the output of the **show route table** command.

The following sections describe how to configure protocol-independent load balancing and how this configuration can affect routing policies:

- [Configuring Load Balancing for Layer 3 VPNs on page 787](#)
- [Configuring Load Balancing and Routing Policies on page 789](#)

### Configuring Load Balancing for Layer 3 VPNs

The configuration of protocol-independent load balancing for Layer 3 VPNs is a little different for IPv4 versus IPv6:

- IPv4—You only need to configure the **multipath** statement at either the **[edit routing-instances *routing-instance-name* routing-options]** hierarchy level or the **[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*]** hierarchy level.
- IPv6—You need to configure the **multipath** statement at both the **[edit routing-instances *routing-instance-name* routing-options]** hierarchy level and the **[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*]** hierarchy level.



**NOTE:** You cannot configure the **multipath** statement and sub-statements at the same time that you have configured the **l3vpn** statement.

To configure protocol-independent load balancing for Layer 3 VPNs, include the **multipath** statement:

```
multipath {
  vpn-unequal-cost equal-external-internal;
}
```

When you include the **multipath** statement at the following hierarchy levels, protocol-independent load balancing is applied to the default routing table for that routing instance (*routing-instance-name.inet.0*):

- **[edit routing-instances *routing-instance-name* routing-options]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]**



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

When you include the **multipath** statement at the following hierarchy levels, protocol-independent load balancing is applied to the specified routing table:

- [edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*]



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

The **vpn-unequal-cost** statement is optional:

- When you include it, protocol-independent load balancing is applied to VPN routes that are equal until the IGP metric with regard to route selection.
- When you do not include it, protocol-independent load balancing is applied to VPN routes that are equal until the router identifier with regard to route selection.



**NOTE:** The **vpn-unequal-cost** statement is not applicable in ACX Series routers.

The **equal-external-internal** statement is also optional. When you include it, protocol-independent load balancing is applied to both internal and external BGP paths. You can configure this in conjunction with egress IP header filtering (enabled with the **vrf-table-label** statement). For more information, see [“Load Balancing and IP Header Filtering for Layer 3 VPNs” on page 770](#).



**NOTE:** You can include the **vpn-unequal-cost equal-external-internal** statement and the **l3vpn** statement at the [edit routing-options forwarding-options chained-composite-next-hop ingress] hierarchy level simultaneously. However, if you do this, EBGp does not work. This means that when there are both paths with chained next hops and paths with nonchained next hops as candidates for EBGp equal-cost multipath (ECMP), the paths using chained next hops are excluded. In a typical case, the excluded paths are the internal paths.



## Configuring Load Balancing and Routing Policies

If you enable protocol-independent load balancing for Layer 3 VPNs by including the **multipath** statement and if you also include the **load-balance per-packet** statement in the routing policy configuration, packets are not load-balanced.

For example, a PE router has the following VRF routing instance configured:

```
[edit routing-instances]
load-balance-example {
  instance-type vrf;
  interface fe-0/1/1.0;
  interface fe-0/1/1.1;
  route-distinguisher 2222:2;
  vrf-target target:2222:2;
  routing-options {
    multipath;
  }
  protocols {
    bgp {
      group group-example {
        import import-policy;
        family inet {
          unicast;
        }
        export export-policy;
        peer-as 4444;
        local-as 3333;
        multipath;
        as-override;
        neighbor 10.12.33.22;
      }
    }
  }
}
```

The PE router also has the following policy statement configured:

```
[edit policy-options policy-statement export-policy]
from protocol bgp;
then {
  load-balance per-packet;
}
```

When you include the **multipath** statement in the VRF routing instance configuration, the paths are no longer marked as BGP paths but are instead marked as multipath paths. Packets from the PE router are not load-balanced.

To ensure that VPN load-balancing functions as expected, do not include the **from protocol** statement in the policy statement configuration. The policy statement should be configured as follows:

```
[edit policy-options policy-statement export-policy]
then {
  load-balance per-packet;
}
```

For more information about how to configure per-packet load balancing, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

- Related Documentation**
- [VPN Per-Packet Load Balancing on page 769](#)
  - *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*

# Enabling BGP PIC for Layer 3 VPNs

- [Configuring BGP PIC Edge for MPLS Layer 3 VPNs on page 791](#)
- [Example: Configuring BGP PIC Edge for MPLS Layer 3 VPNs on page 793](#)

## Configuring BGP PIC Edge for MPLS Layer 3 VPNs

---

In an MPLS VPN Layer 3 environment, it is common for customers to multihome their networks to provide link redundancy. Although the interior gateway protocol (IGP) can provide fast convergence, in certain instances, the time to resolve a link failure and provide an alternate route can be time consuming. For example, a provider edge (PE) router might be configured with 200,000 or more IP prefixes, and a PE router failure could affect many of those prefixes.

BGP Prefix-Independent Convergence (PIC) Edge allows you to install a Layer 3 VPN route in the forwarding table as an alternate path, enabling fast failover when a PE router fails or you lose connectivity to a PE router. This already installed path is used until global convergence through the IGP is resolved. Using the alternative VPN route for forwarding until global convergence is complete reduces traffic loss.

BGP PIC Edge supports multiprotocol BGP IPv4 or IPv6 VPN network layer reachability information (NLRI) resolved using any of these IGP protocols:

- OSPF
- IS-IS
- LDP
- RSVP

BGP PIC Edge does not support multicast traffic.

Before you begin:

1. Configure LDP or RSVP.
2. Configure an IGP: either OSPF or IS-IS.

3. Configure a Layer 3 VPN.
4. Configure multiprotocol BGP for either an IPv4 VPN or an IPv6 VPN.

To configure BGP PIC Edge in an MPLS Layer 3 VPN:

1. Enable BGP PIC Edge:

```
[edit routing-instances routing-instance-name routing-options]
user@host# set protect core
```



**NOTE:** The BGP PIC edge feature is supported on ACX Universal Metro routers and on MX Series 3D Universal Edge routers with MPC interfaces.

2. Configure per-packet load balancing:

```
[edit policy-options]
user@host# set policy-statement policy-name then load-balance per-packet
```

3. Apply the per-packet load balancing policy to routes exported from the routing table to the forwarding table:

```
[edit routing-options forwarding-table]
user@host# set export policy-statement-name
```

4. Verify that BGP PIC Edge is working.

From operational mode, enter the **show route extensive** command:

```
user@host> show route 192.0.2.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
 192.0.2.6/24 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kerne1 192.0.2.6/24 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
    Nexthop: Self
    AS path: [2] 3 I
    Communities: target:2:1
    Path 192.0.2.6 from 192.0.2.4 Vector len 4. Val: 0
..
    #Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x93f4010
    Next-hop reference count: 2
..
    Protocol next hop: 192.0.2.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 192.0.2.5
    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

```

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.



#### BEST PRACTICE:

On MX Series 3D Universal Edge Routers with Modular Port Concentrators (MPCs), we strongly recommend that you enable enhanced IP network services.

To enable enhanced IP network services:

```

[edit chassis]
user@host# set network-services enhanced-ip

```

#### Related Documentation

- [Example: Configuring BGP PIC Edge for MPLS Layer 3 VPNs on page 793](#)
- [Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 878](#)
- [Example: Load Balancing BGP Traffic](#)
- [Network Services Mode Overview](#)
- [Firewall Filters and Enhanced Network Services Mode Overview](#)
- [Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers](#)
- [Configuring Enhanced IP Network Services for a Virtual Chassis](#)

## Example: Configuring BGP PIC Edge for MPLS Layer 3 VPNs

This example shows how to configure BGP prefix-independent convergence (PIC) edge, which allows you to install a Layer 3 VPN route in the forwarding table as an alternate path. This enables fast failover when a provider edge (PE) router fails or you lose connectivity to a PE router. This already installed path is used until global convergence through the interior gateway protocol (IGP) is resolved. Using the alternative VPN route for forwarding until global convergence is complete reduces traffic loss.

- [Requirements on page 794](#)
- [Overview on page 794](#)

- [Configuration on page 795](#)
- [Verification on page 801](#)

## Requirements

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

This example uses the following hardware and software components:

- One MX Series 3D Universal Edge Routers with MPC interfaces to configure the BGP PIC edge feature.
- Five routers that can be a combination of M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- Junos OS Release 13.2 or later on the device with BGP PIC edge configured.

## Overview

In an MPLS VPN Layer 3 environment, it is common for customers to multihome their networks to provide link redundancy. Although the interior gateway protocol (IGP) can provide fast convergence, in certain instances, the time to resolve a link failure and provide an alternate route can be time consuming. For example, a provider edge (PE) router might be configured with 200,000 or more IP prefixes, and a PE router failure could affect many of those prefixes.

This example shows two customer edge (CE) routers, Device CE1 and Device CE2. Devices PE1, PE2, and PE3 are PE routers. Device P1 is a provider core router. Only Device PE1 has BGP PIC edge configured. The example uses the P1-PE2 link (P-PE) link to simulate the loss of a section of the network.


For testing, the address 172.16.1.5/24 is added as a loopback interface address on Device CE2. The address is announced to Device PE2 and Device PE3 and is relayed by way of internal BGP (IBGP) IBGP to Device PE1. On Device PE1, there are two paths to the 172.16.1.5/24 network. These are the primary and a backup path.

## Topology

---

[Figure 76 on page 795](#) shows the sample network.

Figure 76: BGP PIC Edge Scenario



g041992

“CLI Quick Configuration” on page 795 shows the configuration for all of the devices in Figure 76 on page 795.

The section “Step-by-Step Procedure” on page 797 describes the steps on Device PE1.

## Configuration

<b>CLI Quick Configuration</b>	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 <b>[edit]</b> hierarchy level.
<b>Device CE1</b>	<pre> set interfaces ge-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 ebgp type external set protocols bgp group ebgp export send-direct set protocols bgp group ebgp neighbor 10.0.0.2 set policy-options policy-statement send-direct from protocol direct set policy-options policy-statement send-direct then accept set routing-options autonomous-system 101 </pre>
<b>Device CE2</b>	<pre> set interfaces ge-1/2/4 unit 0 family inet address 10.0.0.42/30 set interfaces ge-1/2/3 unit 0 family inet address 10.0.0.46/30 set interfaces lo0 unit 0 family inet address 192.168.0.8/32 set interfaces lo0 unit 0 family inet address 172.16.1.5/24 set protocols bgp group ebgp type external set protocols bgp group ebgp export send-direct set protocols bgp group ebgp neighbor 10.0.0.45 set protocols bgp group ebgp neighbor 10.0.0.41 set policy-options policy-statement send-direct from protocol direct set policy-options policy-statement send-direct then accept set routing-options autonomous-system 102 </pre>
<b>Device P1</b>	<pre> set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.5/30 set interfaces ge-1/2/1 unit 0 family mpls set interfaces ge-1/2/5 unit 0 family inet address 10.0.0.17/30 set interfaces ge-1/2/5 unit 0 family mpls </pre>

```

set interfaces ge-1/2/2 unit 0 family inet address 10.0.0.33/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ge-1/2/5.0
set protocols mpls interface ge-1/2/2.0
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ospf area 0.0.0.0 interface ge-1/2/5.0
set protocols ospf area 0.0.0.0 interface ge-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/1.0
set protocols ldp interface ge-1/2/5.0
set protocols ldp interface ge-1/2/2.0
set protocols ldp interface lo0.0
set routing-options autonomous-system 100

```

**Device PE1**

```

set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 0 family inet address 10.0.0.6/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set protocols mpls interface ge-1/2/1.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.2
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp neighbor 192.168.0.7
set protocols bgp group ibgp neighbor 192.168.0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement lb then load-balance per-packet
set policy-options policy-statement nhs then next-hop self
set routing-instances customer1 instance-type vrf
set routing-instances customer1 interface ge-1/2/0.0
set routing-instances customer1 route-distinguisher 100:1
set routing-instances customer1 vrf-target target:100:1
set routing-instances customer1 routing-options protect core
set routing-instances customer1 protocols bgp group ebgp type external
set routing-instances customer1 protocols bgp group ebgp neighbor 10.0.0.1
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 100
set routing-options forwarding-table export lb

```

**Device PE2**

```

set interfaces ge-1/2/2 unit 0 family inet address 10.0.0.34/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces ge-1/2/3 unit 0 family inet address 10.0.0.45/30
set interfaces lo0 unit 0 family inet address 192.168.0.7/32
set protocols mpls interface ge-1/2/2.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.7
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast

```



```

set protocols bgp group ibgp export nhs
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp neighbor 192.168.0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/2.0
set protocols ldp interface lo0.0
set routing-instances customer1 instance-type vrf
set routing-instances customer1 interface ge-1/2/3.0
set routing-instances customer1 route-distinguisher 100:1
set routing-instances customer1 vrf-target target:100:1
set routing-instances customer1 protocols bgp group ebgp type external
set routing-instances customer1 protocols bgp group ebgp neighbor 10.0.0.46
set routing-options autonomous-system 100

```

**Device PE3**

```

set interfaces ge-1/2/5 unit 0 family inet address 10.0.0.18/30
set interfaces ge-1/2/5 unit 0 family mpls
set interfaces ge-1/2/4 unit 0 family inet address 10.0.0.41/30
set interfaces ge-1/2/4 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set protocols mpls interface ge-1/2/5.0
set protocols mpls interface ge-1/2/4.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.6
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp neighbor 192.168.0.7
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols ospf area 0.0.0.0 interface ge-1/2/5.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/5.0
set protocols ldp interface lo0.0
set routing-instances customer1 instance-type vrf
set routing-instances customer1 interface ge-1/2/4.0
set routing-instances customer1 route-distinguisher 100:1
set routing-instances customer1 vrf-target target:100:1
set routing-instances customer1 protocols bgp group ebgp type external
set routing-instances customer1 protocols bgp group ebgp neighbor 10.0.0.42
set routing-options autonomous-system 100

```

**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 R1:

1. Configure the device interfaces.

```
[edit interfaces]
```

```
user@PE1# set ge-1/2/0 unit 0 family inet address 10.0.0.2/30
```

```
user@PE1# set ge-1/2/1 unit 0 family inet address 10.0.0.6/30
```

```
user@PE1# set ge-1/2/1 unit 0 family mpls
```

```
user@PE1# set lo0 unit 0 family inet address 192.168.0.2/32
```

2. Configure MPLS and LDP on the core-facing interfaces.

```
[edit protocols]
```

```
user@PE1# set mpls interface ge-1/2/1.0
```

```
user@PE1# set ldp interface ge-1/2/1.0
```

```
user@PE1# set ldp interface lo0.0
```

3. Configure an IGP on the core-facing interfaces.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@PE1# set interface ge-1/2/1.0
```

```
user@PE1# set interface lo0.0 passive
```

4. Configure IBGP connections with the other PE devices.

```
[edit protocols bgp group ibgp]
```

```
user@PE1# set type internal
```

```
user@PE1# set local-address 192.168.0.2
```

```
user@PE1# set family inet unicast
```

```
user@PE1# set family inet-vpn unicast
```

```
user@PE1# set export nhs
```

```
user@PE1# set neighbor 192.168.0.7
```

```
user@PE1# set neighbor 192.168.0.6
```

5. Configure the load-balancing policy.

```
[edit policy-options policy-statement lb]
```

```
user@PE1# set then load-balance per-packet
```

6. (Optional) Configure a next-hop self policy.

```
[edit policy-options policy-statement nhs]
```

```
user@PE1# set then next-hop self
```

7. Configure the routing-instance to create the CE-PE EBGp connection.

```
[edit routing-instances customer1]
```

```
user@PE1# set instance-type vrf
```

```
user@PE1# set interface ge-1/2/0.0
```

```
user@PE1# set route-distinguisher 100:1
```

```
user@PE1# set vrf-target target:100:1
```

```
user@PE1# set protocols bgp group ebgp type external
```

```
user@PE1# set protocols bgp group ebgp neighbor 10.0.0.1
```

8. Enable the BGP PIC edge feature.

```
[edit routing-instances customer1]
```

```
user@PE1# set routing-options protect core
```

9. Apply the load-balancing policy.

```
[edit routing-options forwarding-table]
user@PE1# set export lb
```

10. Assign the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set router-id 192.168.0.2
user@PE1# set autonomous-system 100
```

**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 instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
ge-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.6/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.2/32;
    }
  }
}

user@PE1# show protocols
mpls {
  interface ge-1/2/1.0;
}
bgp {
  group ibgp {
    type internal;
    local-address 192.168.0.2;
    family inet {
      unicast;
    }
    family inet-vpn {
      unicast;
    }
    export nhs;
    neighbor 192.168.0.7;
```

```
        neighbor 192.168.0.6;
    }
}
ospf {
    area 0.0.0.0 {
        interface ge-1/2/1.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface ge-1/2/1.0;
    interface lo0.0;
}

user@PE1# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}
policy-statement nhs {
    then {
        next-hop self;
    }
}

user@PE1# show routing-instances
customer1 {
    instance-type vrf;
    interface ge-1/2/0.0;
    route-distinguisher 100:1;
    vrf-target target:100:1;
    routing-options {
        protect core;
    }
    protocols {
        bgp {
            group ebgp {
                type external;
                peer-as 101;
                neighbor 10.0.0.1;
            }
        }
    }
}

user@PE1# show routing-options
router-id 192.168.0.2;
autonomous-system 100;
forwarding-table {
    export lb;
}
```

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

## Verification

Confirm that the configuration is working properly.

- [Displaying Extensive Route Information on page 801](#)
- [Displaying the Forwarding Table on page 803](#)
- [Displaying the OSPF Routes on page 804](#)

### Displaying Extensive Route Information

**Purpose** Confirm that BGP PIC Edge is working.

**Action** From Device PE1, run the `show route extensive table customer1.inet.0 172.16.1/24` command.

```
user@PE1> show route extensive table customer1.inet.0 172.16.1/24

customer1.inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
172.16.1.0/24 (3 entries, 2 announced)
    State: <CalcForwarding>
TSI:
KRT in-kernel 172.16.1.0/24 -> {indirect(262146), indirect(262142)}
Page 0 idx 0, (group ebgp type External) Type 1 val 0x950a62c (adv_entry)
    Advertised metrics:
        Nexthop: Self
        AS path: [100] 102 I
        Communities: target:100:1
Path 172.16.1.0 from 192.168.0.6 Vector len 4. Val: 0
    @BGP    Preference: 170/-101
            Route Distinguisher: 100:1
            Next hop type: Indirect
            Address: 0x9514a74
            Next-hop reference count: 7
            Source: 192.168.0.6
            Next hop type: Router, Next hop index: 990
            Next hop: 10.0.0.5 via ge-1/2/1.0, selected
            Label operation: Push 299824, Push 299856(top)
            Label TTL action: prop-ttl, prop-ttl(top)
            Load balance label: Label 299824: None; Label 299856: None;
            Session Id: 0x280002
            Protocol next hop: 192.168.0.6
            Label operation: Push 299824
            Label TTL action: prop-ttl
            Load balance label: Label 299824: None;
            Indirect next hop: 0x96bc104 262146 INH Session ID: 0x280006
            State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
            Local AS: 100 Peer AS: 100
            Age: 1:38:13 Metric2: 1
            Validation State: unverified
            Task: BGP_100.192.168.0.6+45824
            Announcement bits (1): 1-BGP_RT_Background
            AS path: 102 I
            Communities: target:100:1
            Import Accepted
            VPN Label: 299824
            Localpref: 100
            Router ID: 192.168.0.6
            Primary Routing Table bgp.l3vpn.0
```

```

Indirect next hops: 1
  Protocol next hop: 192.168.0.6 Metric: 1
  Label operation: Push 299824
  Label TTL action: prop-ttl
  Load balance label: Label 299824: None;
  Indirect next hop: 0x96bc104 262146 INH Session ID:

0x280006

  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.0.0.5 via ge-1/2/1.0
    Session Id: 0x280002
  192.168.0.6/32 Originating RIB: inet.3
    Metric: 1                      Node path count: 1
    Forwarding nexthops: 1
      Nexthop: 10.0.0.5 via ge-1/2/1.0

BGP Preference: 170/-101
Route Distinguisher: 100:1
Next hop type: Indirect
Address: 0x9515570
Next-hop reference count: 7
Source: 192.168.0.7
Next hop type: Router, Next hop index: 933
Next hop: 10.0.0.5 via ge-1/2/1.0, selected
Label operation: Push 299856, Push 299872(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 299856: None; Label 299872: None;
Session Id: 0x280002
Protocol next hop: 192.168.0.7
Label operation: Push 299856
Label TTL action: prop-ttl
Load balance label: Label 299856: None;
Indirect next hop: 0x96bc000 262142 INH Session ID: 0x280005
State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

Inactive reason: Not Best in its group - Router ID
Local AS: 100 Peer AS: 100
Age: 1:38:13 Metric2: 1
Validation State: unverified
Task: BGP_100.192.168.0.7+10985
AS path: 102 I
Communities: target:100:1
Import Accepted
VPN Label: 299856
Localpref: 100
Router ID: 192.168.0.7
Primary Routing Table bgp.13vpn.0
Indirect next hops: 1
  Protocol next hop: 192.168.0.7 Metric: 1
  Label operation: Push 299856
  Label TTL action: prop-ttl
  Load balance label: Label 299856: None;
  Indirect next hop: 0x96bc000 262142 INH Session ID:

0x280005

  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.0.0.5 via ge-1/2/1.0
    Session Id: 0x280002
  192.168.0.7/32 Originating RIB: inet.3
    Metric: 1                      Node path count: 1
    Forwarding nexthops: 1
      Nexthop: 10.0.0.5 via ge-1/2/1.0

```

```

#Multipath Preference: 255
  Next hop type: Indirect
  Address: 0x9578010
  Next-hop reference count: 4
  Next hop type: Router, Next hop index: 990
  Next hop: 10.0.0.5 via ge-1/2/1.0, selected
  Label operation: Push 299824, Push 299856(top)
  Label TTL action: prop-ttl, prop-ttl(top)
  Load balance label: Label 299824: None; Label 299856: None;
  Session Id: 0x280002
  Next hop type: Router, Next hop index: 933
  Next hop: 10.0.0.5 via ge-1/2/1.0
  Label operation: Push 299856, Push 299872(top)
  Label TTL action: prop-ttl, prop-ttl(top)
  Load balance label: Label 299856: None; Label 299872: None;
  Session Id: 0x280002
  Protocol next hop: 192.168.0.6
  Label operation: Push 299824
  Label TTL action: prop-ttl
  Load balance label: Label 299824: None;
  Indirect next hop: 0x96bc104 262146 INH Session ID: 0x280006

Weight 0x1
  Protocol next hop: 192.168.0.7
  Label operation: Push 299856
  Label TTL action: prop-ttl
  Load balance label: Label 299856: None;
  Indirect next hop: 0x96bc000 262142 INH Session ID: 0x280005

Weight 0x4000
  State: <ForwardingOnly Int Ext>
  Inactive reason: Forwarding use only
  Age: 1:38:13   Metric2: 1
  Validation State: unverified
  Task: RT
  Announcement bits (1): 0-KRT
  AS path: 102 I
  Communities: target:100:1

```

**Meaning** The Indirect next hop output lines that contain 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.

### Displaying the Forwarding Table

**Purpose** Check the forwarding and kernel routing-table state by using **show route forwarding-table**.

**Action** From Device PE1, run the **show route forwarding-table table customer1 destination 172.16.1.0/24** command.

```
user@PE1> show route forwarding-table table customer1 destination 172.16.1.0/24
```

```

Routing table: customer1.inet
Internet:
Destination      Type RtRef Next hop      Type Index  NhRef Netif
172.16.1.0/24    user  0          10.0.0.5      ulst  262147    2
                                     indr  262146    3
                                     Push 299824, Push 299856(top)
    990          2 ge-1/2/1.0          10.0.0.5      indr  262144    3
                                     Push 300080, Push 299920(top)
    1000          2 ge-1/2/1.0

```

**Meaning** in addition to the forwarding and kernel routing-table state, this command shows the unilist index (262147) used by the Packet Forwarding Engine.

### Displaying the OSPF Routes

**Purpose** Show the OSPF route state.

**Action** From Device PE1, run the **show (ospf | ospf3) route detail** command.

```
user@PE1> show ospf route detail
```

```
betsy@tp0:PE1> show ospf route detail
Topology default Route Table:
```

Prefix	Path Type	Route Type	NH Type	Metric	NextHop Interface	Nexthop Address/LSP
192.168.0.3	Intra	Router	IP	1	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.3, optional-capability 0x0						
192.168.0.6	Intra	Router	IP	2	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.6, optional-capability 0x0						
192.168.0.7	Intra	Router	IP	2	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.7, optional-capability 0x0						
10.0.0.4/30	Intra	Network	IP	1	ge-1/2/1.0	
area 0.0.0.0, origin 192.168.0.3, priority low						
10.0.0.16/30	Intra	Network	IP	2	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.6, priority medium						
10.0.0.32/30	Intra	Network	IP	2	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.7, priority medium						
192.168.0.2/32	Intra	Network	IP	0	lo0.0	
area 0.0.0.0, origin 192.168.0.2, priority low						
192.168.0.3/32	Intra	Network	IP	1	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.3, priority medium						
<b>192.168.0.6/32</b>	Intra	Network	IP	2	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.6, priority medium						
<b>session-id: 2621446, version: 1</b>						
<b>192.168.0.7/32</b>	Intra	Network	IP	2	ge-1/2/1.0	10.0.0.5
area 0.0.0.0, origin 192.168.0.7, priority medium						
<b>session-id: 2621450, version: 1</b>						

**Meaning** The output shows the tracked session IDs for the loopback interface addresses on Devices PE2 and PE3.



- Related Documentation**
- [Configuring BGP PIC Edge for MPLS Layer 3 VPNs on page 791](#)
  - [Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 878](#)



## PART 8

# Configuring Protection and Performance Features for Layer 3 VPNs

- [Configuring Egress Protection in Layer 3 VPNs on page 809](#)
- [Configuring Provider Edge Link Protection in Layer 3 VPNs on page 877](#)
- [Improving Layer 3 VPN Performance on page 911](#)



# Configuring Egress Protection in Layer 3 VPNs

- [Egress Protection for BGP Labeled Unicast on page 809](#)
- [Configuring Egress Protection for BGP Labeled Unicast on page 811](#)
- [Example: Configuring Egress Protection for BGP Labeled Unicast on page 813](#)
- [Example: Configuring MPLS Egress Protection for Layer 3 VPN Services on page 826](#)
- [Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP on page 843](#)

## Egress Protection for BGP Labeled Unicast

---


When network node or link failures occur, it takes some time to restore service using traditional routing table convergence. Local repair procedures can provide much faster restoration by establishing local protection as close to a failure as possible. Fast protection for egress nodes is available to services in which BGP labeled unicast interconnects IGP areas, levels, or autonomous systems (ASs). If a provider router detects that an egress router (AS or area border router) is down, it immediately forwards the traffic destined to that router to a protector router that forwards the traffic downstream to the destination.

To provide egress protection for BGP labeled unicast, the protector node must create a backup state for downstream destinations before the failure happens. The basic idea of the solution is that the protector node constructs a forwarding state associated with the protected node and relays the MPLS labels assigned by the protected node further downstream to the final destination.

This feature supports the applications Inter-AS Option C and Seamless MPLS.

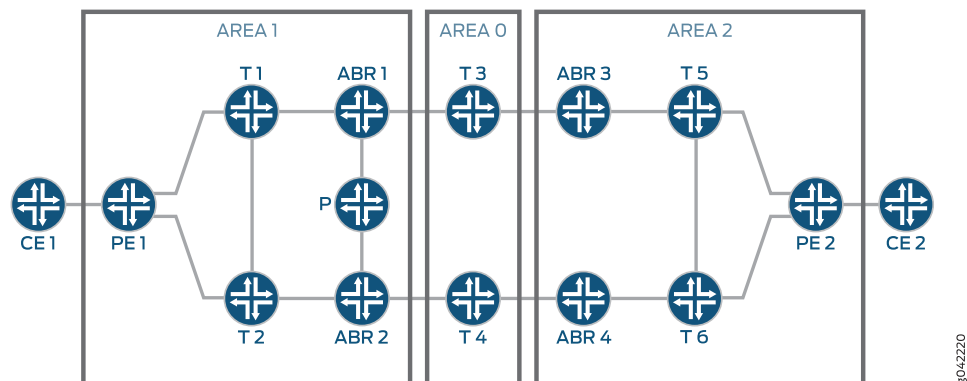
*Inter-AS Option C*—BGP labeled unicast provides end-to-end transport label-switched paths (LSPs) by stitching the intra-AS LSPs together. AS boundary routers run EBGp to other AS boundary routers to exchange labels for /32 PE loopback routes. IBGP runs between the provider edge router and AS boundary routers within each AS. In [Figure 77 on page 810](#), the traffic goes from CE1 to CE2. ASBR1 is the protected AS boundary router, ASBR2 is the protector, and Device P1 is the point of local repair (PLR). The primary path is chosen from PE1 to PE2 over ASBR1 and ASBR3. When ASBR1 fails, Router P1 detects the ASBR1 failure and forwards the traffic to ASBR2, which provides backup service and forwards the traffic downstream.

Figure 77: Inter-AS Option C



*Seamless MPLS*—BGP labeled unicast provides end-to-end transport LSPs by stitching the intra-area/level LSPs. Area border routers (ABRs) run BGP labeled unicast to other ABRs to exchange labels for /32 PE loopback routes. In Figure 78 on page 810, the traffic goes from Device CE1 to Device CE2. ABR1 is the protected ABR, ABR2 is the protector, and T1 is the PLR. The primary path is chosen from PE1 to PE2 over ABR1 and ABR3. When ABR1 fails, Router T1 detects the ABR1 failure and forwards the traffic to ABR2, which provides backup service and forwards the traffic downstream.

Figure 78: Seamless MPLS



In each of these applications, the protected node advertises a primary BGP labeled unicast route that needs protection. When fast protection is enabled, BGP advertises the label routes with a special address as the next hop. This special address is a context identifier that is configured through the CLI. The protected node also advertises the context identifier in IGP and a NULL label in LDP for the context identifier.

The backup node advertises backup BGP labeled unicast routes for the protected routes. The protector node forwards traffic to the backup node using the labels advertised by the backup node.

The protector node provides the backup service by cross-connecting the labels originated by the protected node and the labels originated by the backup node. The protector node forwards the traffic to the backup node in case of failure of the protected node. The protector node advertises the same context-identifier into IGP with high metric. Also, it advertises a real label in LDP for the context identifier. The protector node listens for the BGP labeled unicast routes advertised by both the protected node and backup node and populates the context label table and backup FIB. When traffic with the real context LDP

label arrives, the lookup is done in the context of a protected node. The protector node often acts as the backup node.

The PLR detects the protected node failure and forwards the MPLS traffic to the protector node. The high IGP metric along with the LDP label advertised by the protector node ensure that the PLR uses the protector node as an LDP backup LSP.

There are two supported protection types: collocated protector and centralized protector. In the collocated type, the protector node is also the backup node. In the centralized type, the backup node is different from the protector node.

#### Related Documentation

- [Configuring Egress Protection for BGP Labeled Unicast on page 811](#)
- [Example: Configuring Egress Protection for BGP Labeled Unicast on page 813](#)
- [egress-protection on page 1052](#)

## Configuring Egress Protection for BGP Labeled Unicast

Fast protection for egress nodes is available to services in which BGP labeled unicast interconnects IGP areas, levels, or ASs. If a provider router detects that an egress router (AS or area border router) is down, it immediately forwards the traffic destined to that router to a protector router that forwards the traffic downstream to the destination.

Before configuring egress protection for BGP labeled unicast, ensure that all routers in the AS or area are running Junos OS 14.1 or a later release.

To configure egress protection for BGP labeled unicast:

1. Add the following configuration to the *protected* router:

```
[edit protocols]
  mpls {
    egress-protection {
      context-identifier context-id {
        primary;
      }
    }
  }
  bgp {
    group group-name {
      type internal;
      family inet {
        labeled-unicast {
          egress-protection {
            context-identifier context-id;
          }
        }
      }
    }
  }
}
```

2. Add the following configuration to the *protector* router:

```
[edit protocols]
  mpls {
    egress-protection {
      context-identifier context-id {
        protector;
      }
    }
  }
  bgp {
    group group-name {
      type internal;
      family inet {
        labeled-unicast {
          egress-protection;
        }
      }
    }
  }
}
```

3. Add the following configuration to the *PLR* (point of local repair) router:

```
[edit protocols]
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  isis {
    backup-spf-options per-prefix-calculation;
    level 1 disable;
    interface all {
      node-link-protection;
    }
  }
  ldp {
    track-igp-metric;
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

4. Run **show bgp neighbor** on the protected router to verify that egress protection is enabled, for example:

```
user@host# run show bgp neighbor
Peer: 192.0.2.2+179 AS 65536 Local: 192.0.2.1+59264 AS 65536
Type: Internal State: Established Flags: <ImportEval Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress KeepAll AddressFamily Rib-group Refresh>
Address families configured: inet-label-unicast
Local Address: 192.0.2.1 Holdtime: 90 Preference: 170
NLRI configured with egress-protection: inet-label-unicast
Egress-protection NLRI inet-label-unicast
Number of flaps: 0
```



- Related Documentation**
- [Egress Protection for BGP Labeled Unicast on page 809](#)
  - [Example: Configuring Egress Protection for BGP Labeled Unicast on page 813](#)
  - [egress-protection \(MPLS\) on page 1053](#)

---

## Example: Configuring Egress Protection for BGP Labeled Unicast

---

This example shows how to configure BGP labeled unicast protection that can be used in case of a PE failure in an Inter-AS Option C topology.

- [Requirements on page 813](#)
- [Overview on page 813](#)
- [Configuration on page 814](#)
- [Verification on page 825](#)

### Requirements

This example uses the following hardware and software components:

- M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers
- Junos OS Release 14.1 or later

### Overview

When network node or link failures occur, it takes some time to restore service using traditional routing table convergence. Local repair procedures can provide much faster restoration by establishing local protection as close to a failure as possible. Fast protection for egress nodes is available to services in which BGP labeled unicast interconnects IGP areas, levels, or autonomous systems (ASs). If a provider router detects that an egress router (AS or area border router) is down, it immediately forwards the traffic destined to that router to a protector router that forwards the traffic downstream to the destination.

This example shows how to configure labeled-unicast egress protection in a Layer 3 VPN.

---

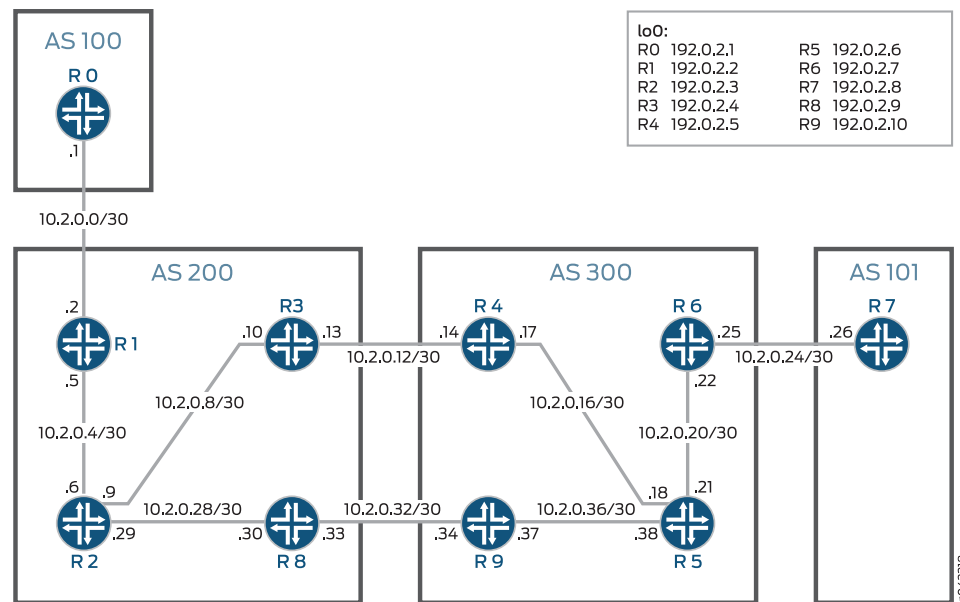
### Topology

---

In this example, an Inter-AS Option C topology is set up by configuring two customer edge (CE) devices and six service provider edge (PE) devices in four autonomous systems. The CE devices are configured in AS100 and AS101. The PE devices are configured in AS200 and AS300.

[Figure 79 on page 814](#) shows the topology used in this example.

Figure 79: Egress Protection in a Layer 3 VPN



The aim of this example is to protect PE Router R4. Egress protection is configured on Router R4 and Router R9 so that the traffic can be routed through the backup link (R9 to R8) when Router R4 (or the link from R5 to R4) goes down. In this example, Router R4 is the protected router, Router R9 is the protector router, and Router R5 is the point of local repair (PLR).

## 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.

**Router R0**

```

set interfaces ge-0/0/0 unit 0 description toR1
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.1/30
set interfaces lo0 unit 0 family inet address 192.0.2.1/24 primary
set routing-options router-id 192.0.2.1
set routing-options autonomous-system 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10

```

**Router R1**

```

set interfaces ge-0/0/0 unit 0 description toR0
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR2
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.5/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.2/24
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 200
set protocols mpls label-switched-path ToR3 to 192.0.2.4

```

```

set protocols mpls label-switched-path ToR8 to 192.0.2.9
set protocols mpls interface all
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.2
set protocols bgp group parent-vpn-peers family inet unicast
set protocols bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
set protocols bgp group parent-vpn-peers neighbor 192.0.2.4
set protocols bgp group parent-vpn-peers neighbor 192.0.2.9
set protocols bgp group toR6 type external
set protocols bgp group toR6 multihop ttl 10
set protocols bgp group toR6 local-address 192.0.2.2
set protocols bgp group toR6 family inet-vpn unicast
set protocols bgp group toR6 peer-as 300
set protocols bgp group toR6 neighbor 192.0.2.7
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 metric 10
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement child_vpn_routes term 1 from protocol bgp
set policy-options policy-statement child_vpn_routes term 1 then accept
set policy-options policy-statement child_vpn_routes term 2 then reject
set policy-options policy-statement vpnexport term 1 from protocol ospf
set policy-options policy-statement vpnexport term 1 then community add test_comm
set policy-options policy-statement vpnexport term 1 then accept
set policy-options policy-statement vpnexport term 2 then reject
set policy-options policy-statement vpnimport term 1 from protocol bgp
set policy-options policy-statement vpnimport term 1 from community test_comm
set policy-options policy-statement vpnimport term 1 then accept
set policy-options policy-statement vpnimport term 2 then reject
set policy-options community text_comm members target:1:200
set routing-instances customer-provider-vpn instance-type vrf
set routing-instances customer-provider-vpn interface ge-0/0/0.0
set routing-instances customer-provider-vpn route-distinguisher 192.0.2.4:1
set routing-instances customer-provider-vpn vrf-import vpnimport
set routing-instances customer-provider-vpn vrf-export vpnexport
set routing-instances customer-provider-vpn vrf-target target:200:1
set routing-instances customer-provider-vpn protocols ospf export child_vpn_routes
set routing-instances customer-provider-vpn protocols ospf area 0.0.0.0 interface
  ge-0/0/0.0

```

**Router R2**

```

set interfaces ge-0/0/0 unit 0 description toR3
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.9/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR1
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.6/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 description toR8
set interfaces ge-0/0/2 unit 0 family inet address 10.2.0.29/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.3/24
set routing-options router-id 192.0.2.3
set routing-options autonomous-system 200
set protocols mpls interface all
set protocols ospf traffic-engineering

```

```

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 metric 10
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface lo0.0

```

**Router R3**

```

set interfaces ge-0/0/0 unit 0 description toR2
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.10/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR4
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.13/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.4/24
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 200
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path ToR1 to 192.0.2.2
set protocols mpls interface all
set protocols bgp group toR4 type external
set protocols bgp group toR4 family inet unicast
set protocols bgp group toR4 family inet labeled-unicast rib inet.3
set protocols bgp group toR4 export send-pe
set protocols bgp group toR4 neighbor 10.2.0.14 peer-as 300
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.4
set protocols bgp group parent-vpn-peers family inet unicast
set protocols bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
set protocols bgp group parent-vpn-peers export next-hop-self
set protocols bgp group parent-vpn-peers neighbor 192.0.2.2
set protocols bgp group parent-vpn-peers neighbor 192.0.2.9
set protocols ospf traffic-engineering
set protocols ospf export from-bgp
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self term 1 then next-hop-self
set policy-options policy-statement send-pe from route-filter 192.0.2.2/24 exact
set policy-options policy-statement send-pe then accept

```

**Router R4**

```

set interfaces ge-0/0/0 unit 0 description toR5
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.17/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR3
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.14/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.5/24
set interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.5507.2049.00

```

```

set routing-options router-id 192.0.2.5
set routing-options autonomous-system 300
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path ToR6 to 192.0.2.7
set protocols mpls interface all
set protocols mpls interface fxp.0 disable
set protocols mpls egress-protection context-identifier 203.0.113.1 primary
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.5
set protocols bgp group parent-vpn-peers family inet unicast
set protocols bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
set protocols bgp group parent-vpn-peers family inet labeled-unicast egress-protection
context-identifier 203.0.113.1
set protocols bgp group parent-vpn-peers export next-hop-self
set protocols bgp group parent-vpn-peers neighbor 192.0.2.7
set protocols bgp group parent-vpn-peers neighbor 192.0.2.10
set protocols bgp group toR3 type external
set protocols bgp group toR3 family inet labeled-unicast rib inet.3
set protocols bgp group toR3 export send-pe
set protocols bgp group toR3 peer-as 200
set protocols bgp group toR3 neighbor 10.2.0.13
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface ge-0/0/0.0 level 2 metric 10
set protocols isis interface lo0.0 passive
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self term 1 then next-hop-self
set policy-options policy-statement send-pe from route-filter 192.0.2.7/24 exact
set policy-options policy-statement send-pe then accept

```

**Router R5**

```

set interfaces ge-0/0/0 unit 0 description toR4
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.18/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR6
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.21/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 description toR9
set interfaces ge-0/0/2 unit 0 family inet address 10.2.0.38/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.6/24
set interfaces lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5507.2050.00
set routing-options router-id 192.0.2.6
set routing-options autonomous-system 300
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis backup-spf-options per-prefix-calculation
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all node-link-protection

```

```

set protocols isis interface fxp0.0 disable
set protocols isis interface ge-0/0/0.0 link-protection
set protocols isis interface ge-0/0/0.0 level 2 metric 10
set protocols isis interface ge-0/0/1.0 link-protection
set protocols isis interface ge-0/0/1.0 level 2 metric 10
set protocols isis interface ge-0/0/2.0 link-protection
set protocols isis interface ge-0/0/2.0 level 2 metric 10
set protocols isis interface lo0.0 passive
set protocols ldp track-igp-metric
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable

```

```

Router R6
set interfaces ge-0/0/0 unit 0 description toR7
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.25/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR5
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.22/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.7/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2048.00
set routing-options router-id 192.0.2.7
set routing-options autonomous-system 300
set protocols mpls label-switched-path ToR4 to 192.0.2.5
set protocols mpls label-switched-path ToR9 to 192.0.2.10
set protocols mpls interface all
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.7
set protocols bgp group parent-vpn-peers family inet unicast
set protocols bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
set protocols bgp group parent-vpn-peers neighbor 192.0.2.5
set protocols bgp group parent-vpn-peers neighbor 192.0.2.10
set protocols bgp group toR1 type external
set protocols bgp group toR1 multihop ttl 10
set protocols bgp group toR1 local-address 192.0.2.7
set protocols bgp group toR1 family inet-vpn unicast
set protocols bgp group toR1 peers-as 200
set protocols bgp group toR1 neighbor 192.0.2.2
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface ge-0/0/1.0 level 2 metric 10
set protocols isis interface lo0.0 passive
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement child-vpn-routes term 1 from protocol bgp
set policy-options policy-statement child-vpn-routes term 1 then accept
set policy-options policy-statement child-vpn-routes term 2 then reject
set policy-options policy-statement vpnexport term 1 from protocol ospf
set policy-options policy-statement vpnexport term 1 then community add test_comm
set policy-options policy-statement vpnexport term 1 then accept
set policy-options policy-statement vpnexport term 2 then reject
set policy-options policy-statement vpnimport term 1 from protocol bgp
set policy-options policy-statement vpnimport term 1 from community test_comm

```

```

set policy-options policy-statement vpnimport term 1 then accept
set policy-options policy-statement vpnimport term 2 then reject
set policy-options community test_comm members target:1:300
set routing-instances customer-provider-vpn instance-type vrf
set routing-instances customer-provider-vpn interface ge-0/0/0.0
set routing-instances customer-provider-vpn route-distinguisher 192.0.2.5:1
set routing-instances customer-provider-vpn vrf-import vpnimport
set routing-instances customer-provider-vpn vrf-export vpnexport
set routing-instances customer-provider-vpn vrf-target target:300:1
set routing-instances customer-provider-vpn protocols ospf export child-vpn-routes
set routing-instances customer-provider-vpn protocols ospf area 0.0.0.0 interface
  ge-0/0/0.0

```

**Router R7**

```

set interfaces ge-0/0/0 unit 0 description toR6
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.26/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.8/24 primary
set routing-options router-id 192.0.2.8
set routing-options autonomous-system 101
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 10

```

**Router R8**

```

set interfaces ge-0/0/0 unit 0 description toR9
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.33/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR2
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.30/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.9/24
set routing-options router-id 192.0.2.9
set routing-options autonomous-system 200
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path ToR1 to 192.0.2.2
set protocols mpls interface all
set protocols bgp group toR9 type external
set protocols bgp group toR9 family inet unicast
set protocols bgp group toR9 family inet labeled-unicast rib inet.3
set protocols bgp group toR9 export send-pe
set protocols bgp group toR9 neighbor 10.2.0.34 peer-as 300
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.9
set protocols bgp group parent-vpn-peers family inet unicast
set protocols bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
set protocols bgp group parent-vpn-peers export next-hop-self
set protocols bgp group parent-vpn-peers neighbor 192.0.2.2
set protocols bgp group parent-vpn-peers neighbor 192.0.2.4
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 metric 10
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement from-bgp from protocol bgp
set policy-options policy-statement from-bgp then metric add 100

```

```

set policy-options policy-statement from-bgp then accept
set policy-options policy-statement next-hop-self term 1 then next-hop-self
set policy-options policy-statement send-pe from route-filter 192.0.2.2/24 exact
set policy-options policy-statement send-pe then accept

```

```

Router R9
set interfaces ge-0/0/0 unit 0 description toR8
set interfaces ge-0/0/0 unit 0 family inet address 10.2.0.34/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 description toR5
set interfaces ge-0/0/1 unit 0 family inet address 10.2.0.37/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.10/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2062.00
set routing-options router-id 192.0.2.10
set routing-options autonomous-system 300
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path ToR6 to 192.0.2.7
set protocols mpls interface all
set protocols mpls egress-protection context-identifier 203.0.113.1 protector
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.10
set protocols bgp group parent-vpn-peers family inet unicast
set protocols bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
set protocols bgp group parent-vpn-peers family inet labeled-unicast egress-protection
set protocols bgp group parent-vpn-peers export next-hop-self
set protocols bgp group parent-vpn-peers neighbor 192.0.2.7
set protocols bgp group parent-vpn-peers neighbor 192.0.2.5
set protocols bgp group toR8 type external
set protocols bgp group toR8 family inet labeled-unicast rib inet.3
set protocols bgp group toR8 export send-pe
set protocols bgp group toR8 neighbor 10.2.0.33 peer-as 200
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface ge-0/0/1.0 level 2 metric 10
set protocols isis interface lo0.0 passive
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self term 1 then next-hop-self
set policy-options policy-statement send-pe from route-filter 192.0.2.7/24 exact
set policy-options policy-statement send-pe then accept

```

### Configuring Egress Protection in Layer 3 VPNs

**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 labeled unicast egress protection:

1. Configure the interfaces on each router, for example:  
[edit interfaces]



```

user@R4# set ge-0/0/0 unit 0 description toR5
user@R4# set ge-0/0/0 unit 0 family inet address 10.2.0.17/30
user@R4# set ge-0/0/0 unit 0 family iso
user@R4# set ge-0/0/0 unit 0 family mpls

user@R4# set ge-0/0/1 unit 0 description toR3
user@R4# set ge-0/0/1 unit 0 family inet address 10.2.0.14/30
user@R4# set ge-0/0/1 unit 0 family mpls

user@R4# set lo0 unit 0 family inet address 192.0.2.5/24
user@R4# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5507.2049.00

```

2. Configure the router ID and autonomous system (AS) number for each router, for example:

```

[edit routing-options]
user@R4# set router-id 192.0.2.5
user@R4# set autonomous-system 300

```

In this example, the router ID is chosen to be identical to the loopback address configured on the router.

3. Configure the protocols on each router, for example:

```

[edit protocols]
user@R4# set mpls traffic-engineering bgp-igp-both-ribs
user@R4# set mpls label-switched-path ToR6 to 192.0.2.7
user@R4# set mpls interface all
user@R4# set mpls interface fxp.0 disable
user@R4# set bgp group parent-vpn-peers type internal
user@R4# set bgp group parent-vpn-peers local-address 192.0.2.5
user@R4# set bgp group parent-vpn-peers family inet unicast
user@R4# set bgp group parent-vpn-peers family inet labeled-unicast rib inet.3
user@R4# set bgp group parent-vpn-peers export next-hop-self
user@R4# set bgp group parent-vpn-peers neighbor 192.0.2.7
user@R4# set bgp group parent-vpn-peers neighbor 192.0.2.10
user@R4# set bgp group toR3 type external
user@R4# set bgp group toR3 family inet labeled-unicast rib inet.3
user@R4# set bgp group toR3 export send-pe
user@R4# set bgp group toR3 peer-as 200
user@R4# set bgp group toR3 neighbor 10.2.0.13
user@R4# set isis level 1 disable
user@R4# set isis level 2 wide-metrics-only
user@R4# set isis interface ge-0/0/0.0 level 2 metric 10
user@R4# set isis interface lo0.0 passive
user@R4# set ldp interface ge-0/0/0.0
user@R4# set ldp interface ge-0/0/1.0
user@R4# set ldp interface lo0.0

```

4. Configure routing policies on all PE routers and AS border routers (Routers R1, R3, R4, R6, R8, and R9), for example:

```

user@R4# set policy-options policy-statement next-hop-self term 1 then
next-hop-self

```

```
user@R4# set policy-options policy-statement send-pe from route-filter 192.0.2.7/24
exact
user@R4# set policy-options policy-statement send-pe then accept
```

5. Configure the VPN routing instance on Routers R1 and R6.

```
user@R1# set routing-instances customer-provider-vpn instance-type vrf
user@R1# set routing-instances customer-provider-vpn interface ge-0/0/0.0
user@R1# set routing-instances customer-provider-vpn route-distinguisher
192.0.2.4:1
user@R1# set routing-instances customer-provider-vpn vrf-import vpnimport
user@R1# set routing-instances customer-provider-vpn vrf-export vpnexport
user@R1# set routing-instances customer-provider-vpn vrf-target target:200:1
user@R1# set routing-instances customer-provider-vpn protocols ospf export
child_vpn_routes
user@R1# set routing-instances customer-provider-vpn protocols ospf area 0.0.0.0
interface ge-0/0/0.0
```

and

```
user@R6# set routing-instances customer-provider-vpn instance-type vrf
user@R6# set routing-instances customer-provider-vpn interface ge-0/0/0.0
user@R6# set routing-instances customer-provider-vpn route-distinguisher
192.0.2.5:1
user@R6# set routing-instances customer-provider-vpn vrf-import vpnimport
user@R6# set routing-instances customer-provider-vpn vrf-export vpnexport
user@R6# set routing-instances customer-provider-vpn vrf-target target:300:1
user@R6# set routing-instances customer-provider-vpn protocols ospf export
child-vpn-routes
user@R6# set routing-instances customer-provider-vpn protocols ospf area 0.0.0.0
interface ge-0/0/0.0
```

6. Configure egress protection for Router R4, setting Router R4 as the protected router and Router R9 as the protector.

```
user@R4# set protocols mpls egress-protection context-identifier 203.0.113.1 primary
user@R4# set protocols bgp group parent-vpn-peers family inet labeled-unicast
egress-protection context-identifier 203.0.113.1
```

and

```
user@R9# set protocols mpls egress-protection context-identifier 203.0.113.1
protector
user@R9# set protocols bgp group parent-vpn-peers family inet labeled-unicast
egress-protection
```

---

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, **show policy-options** (if applicable), and **show routing-instances** (if applicable) commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R4# show interfaces
ge-0/0/0 {
  unit 0 {
    description toR5;
    family inet {
      address 10.2.0.17/30;
    }
    family iso;
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    description toR3;
    family inet {
      address 10.2.0.14/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.5/24;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.5507.2049.00;
    }
  }
}

```

```

user@R4# show routing-options
router-id 192.0.2.5;
autonomous-system 300;

```

```

user@R4# show protocols
mpls {
  traffic-engineering bgp-igp-both-ribs;
  label-switched-path ToR6 {
    to 192.0.2.7;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 203.0.113.1 {
      primary;
    }
  }
}
bgp {
  group parent-vpn-peers {
    type internal;
    local-address 192.0.2.5;
    family inet {

```

```

        unicast;
        labeled-unicast {
            rib {
                inet.3;
            }
            egress-protection {
                context-identifier {
                    203.0.113.1;
                }
            }
        }
    }
    export next-hop-self;
    neighbor 192.0.2.7;
    neighbor 192.0.2.10;
}
group toR3 {
    type external;
    family inet {
        unicast;
        labeled-unicast {
            rib {
                inet.3;
            }
        }
    }
    export send-pe;
    peer-as 200;
    neighbor 10.2.0.13;
}
}
isis {
    level 1 disable;
    level 2 wide-metrics-only;
    interface ge-0/0/0.0 {
        level 2 metric 10;
    }
    interface lo0.0 {
        passive;
    }
}
ldp {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
    interface lo0.0;
}

```

```

user@R4# show policy-options
policy-statement next-hop-self {
    term 1 {
        then {
            next-hop self;
        }
    }
}
policy-statement send-pe {
    from {
        route-filter 192.0.2.7/24 exact;
    }
}

```

```

    then accept;
}

```

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

Repeat the procedure for every router in this example, using the appropriate interface names and addresses for each router.

## Verification

- [Verifying That Egress Protection Is Enabled on page 825](#)
- [Verifying the State of the Protected ASBR as 'primary' on page 825](#)
- [Verifying the State of the Protector ASBR as 'protector' on page 825](#)

### Verifying That Egress Protection Is Enabled

**Purpose** Verify that egress protection is enabled on the protected router, Router R4.

**Action** Run **show bgp neighbor** on Router R4 to verify that egress protection is enabled.

```

user@R4> show bgp neighbor
Peer: 192.0.2.10+45824 AS 300    Local: 192.0.2.5+27630 AS 300
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ next-hop-self ]
  Options: <Preference LocalAddress AddressFamily Refresh>
  Address families configured: inet-unicast inet-labeled-unicast
  Local Address: 192.0.2.5 Holdtime: 90 Preference: 170
  NLRI configured with egress-protection: inet-labeled-unicast
  Egress-protection NLRI inet-labeled-unicast context-identifier: 203.0.113.1
  Number of flaps: 0
  ...

```

### Verifying the State of the Protected ASBR as 'primary'

**Purpose** Verify that the state of the protected AS border router, Router R4, is 'primary'.

**Action** Run **show mpls context-identifier** on Router R4.

```

user@R4> show mpls context-identifier
ID          Type      Metric    ContextTable
203.0.113.1  primary    1
Total 1, Primary 1, Protector 0

```

### Verifying the State of the Protector ASBR as 'protector'

**Purpose** Verify that the state of the protector AS border router, Router R9, is 'protector'.

**Action** Run `show mpls context-identifier` on Router R9.

```
user@R9> show mpls context-identifier
ID          Type      Metric  ContextTable
203.0.113.1 protector 16777215 __203.0.113.1__.mpls.0
Total 1, Primary 0, Protector 1
```

- Related Documentation**
- [Egress Protection for BGP Labeled Unicast on page 809](#)
  - [Configuring Egress Protection for BGP Labeled Unicast on page 811](#)
  - [egress-protection \(BGP\) on page 1052](#)

---

## Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

---

This example describes a local repair mechanism for protecting Layer 3 VPN services against egress provider edge (PE) router failure in a scenario where the customer edge (CE) routers are multihomed with more than one PE router.

The following terminology is used in this example:

- **Originator PE router**—A PE router with protected routing instances or subnets that distributes the primary Layer 3 VPN router.
- **Backup PE router**—A PE router that announces a backup Layer 3 VPN route.
- **Protector PE router**—A router that cross-connects VPN labels distributed by the originator PE router to the labels originated by the backup PE router. The protector PE router can also be a backup PE router.
- **Transport LSP**—An LDP-signaled label-switched path (LSP) for BGP next hops.
- **PLR**—A router acting as the point of local repair (PLR) that can redirect Layer 3 VPN traffic to a protector PE router to enable fast restoration and reroute.
- **Loop-free alternate routes**—A technology that essentially adds IP fast-reroute capability for the interior gateway protocol (IGP) by precomputing backup routes for all the primary routes of the IGP. In the context of this document, the IGP is IS-IS.
- **Multihoming**—A technology that enables you to connect a CE device to multiple PE routers. In the event that a connection to the primary PE router fails, traffic can be automatically switched to the backup PE router.
- **Context identifier**—An IPv4 address used to identify the VPN prefix that requires protection. The identifier is propagated to the PE and PLR core routers, making it possible for the protected egress PE router to signal the egress protection to the protector PE router.
- **Dual protection**—A protection mechanism where two PE routers can simultaneously act as the primary PE router and the protector PE router for their respective context ID routes or next hops. For example, between the two PE routers PE1 and PE2, PE1 could be a primary PE router for context identifier 203.0.113.1 and protector for context

identifier 203.0.113.2 Likewise, the PE2 router could be a protector for context identifier 203.0.113.1 and a primary PE router for context identifier 203.0.113.2.

This example contains the following topics:


- [Egress Protection for Layer 3 VPN Edge Protection Overview on page 827](#)
- [Example: Configuring Egress Protection for Layer 3 VPN Services on page 834](#)

## Egress Protection for Layer 3 VPN Edge Protection Overview

Typically, Layer 3 VPN service restoration for multihomed customer edge (CE) routers depends on the ingress provider edge (PE) router to detect the egress PE link or node failure and switch traffic to the backup PE router. To achieve faster restoration, a protector mechanism for the PE router can be used to perform local restoration of the service immediately in case of an egress PE node failure. This mechanism requires the router at the point of local repair (PLR) to redirect VPN traffic to a protector PE router for fast reroute of traffic.

The following topology describes the concept of egress protection.

*Figure 80: Sample Topology for Egress Protection*



In this topology:

Router PE3 acts as the protector for the PE2 Layer 3 VPN routing instances or subnets.

The CE routers are part of a VPN where Router CE1 is multihomed with Router PE1 and Router PE2. Likewise, Router CE2 is multihomed with Routers PE2 and PE3.

Router PE1 can be the originator for the context identifier for Router CE1, while Router PE2 is the protector for that context identifier. Likewise, PE2 can be the originator for the context identifier for Router CE2, while Router PE3 is the protector for that context identifier.

The working path taken by Router PE4 might be through PLR>PE2 for both Router CE1 and Router CE2. The backup path for Router CE1 is through PLR>PE1. The backup path for Router CE2 is through PLR>PE3. Traffic flows through the working path under normal circumstances.

When Router PE4 detects a PE2 node or link failure, traffic is rerouted from the working path to the protected path. In the normal failover process, the detection of failure and the recovery rely on the control plane and is therefore relatively slow.

Typically, if there is a link or node failure in the core network, the egress PE router would have to rely on the ingress PE router to detect the failure and switch over to the backup path, because a local repair option for egress failure is not available.

To provide a local repair solution for the egress PE link or node failure, a mechanism known as egress protection can be used to repair and restore the connection quickly. If egress protection is configured, the PLR router detects the PE2 link or node failure and reroutes traffic through the protector Router PE3 using the backup LDP-signaled label-switched path (LSP). The PLR router uses per-prefix loop-free alternate routes to program the backup next hop through Router PE3, and traffic is forwarded to Routers CE1 and CE2 using the alternate paths. This restoration is done quickly after the PLR router detects the Router PE2 egress node or link failure.

The dual protection mechanism can also be used for egress protection where the two PE routers can simultaneously act as the primary PE router and the protector PE router for their respective context ID routes or next hops.

---

### Router Functions

In [Figure 80 on page 827](#), the following routers perform the following functions:

#### ***Protected PE Router***

The protected PE, PE2, performs the following functions:

- Updates a context identifier for the BGP next hop for the Layer 3 VPN prefix.
- Advertises the context identifier to the IS-IS domain.

#### ***Protector PE Router***

The protector PE router, PE3, performs the following functions:

- Advertises the context identifier to the IS-IS domain with a high metric. The high IGP metric (configurable) along with the LDP label ensures that the PLR router uses the LDP-signaled backup LSP in the event of an egress PE router failure.
- Builds a context-label table for route lookup and a backup forwarding table for the protected PE router (PE2).



**NOTE:** The protector PE router should not be in the forwarding path to the primary PE router.

---



### PLR Router

The router acting as the point of local repair (PLR) performs the following functions:

- Computes per-prefix loop-free alternate routes. For this computation to work, the configuration of the **node-link-protection** statement and the **backup-spf-options per-prefix-calculation** statement is necessary at the **[edit protocols isis]** hierarchy level.
- Installs backup next hops for the context identifier through the PE3 router (protector PE).
- Detects PE router failure and redirects the transport LSP traffic to the protector.



**NOTE:** The PLR router must be directly connected to the protector router (in this case, PE3). If not, the loop-free alternate route cannot find the backup path to the protector. This limitation is removed in Junos OS Release 13.3 and later.

### Protector and Protection Models

Protector is a new role or function for the restoration of egress PE node failure. This role could be played by a backup egress PE router or any other node that participates in the VPN control plane for VPN prefixes that require egress node protection. There are two protection models based on the location and role of a protector:

- **Co-located protector**—In this model, the protector PE router and the backup PE router configurations are done on the same router. The protector is co-located with the backup PE router for the protected prefix, and it has a direct connection to the multihomed site that originates the protected prefix. In the event of an egress PE failure, the protector receives traffic from the PLR router and routes the traffic to the multihomed site.
- **Centralized protector**—In this model, the protector PE router and the backup PE router are different. The centralized protector might not have a direct connection to the multihomed site. In the event of an egress PE link or node failure, the centralized protector reroutes the traffic to the backup egress PE router with the VPN label advertised for the backup egress PE router that takes over the role of sending traffic to the multihomed site.

A network can use either of the protection models or a combination of both, depending on the requirement.

As a special scenario of egress node protection, if a router is both a Protector and a PLR, it installs backup next hops to protect the transport LSP. In particular, it does not need a bypass LSP for local repair.

In the Co-located protector model, the PLR or the Protector is directly connected to the CE via a backup AC, while in the Centralized protector model, the PLR or the protector has an MPLS tunnel to the backup PE. In either case, the PLR or the Protector will install a backup next hop with a label followed by a lookup in a **context label** table, i.e. **\_\_context\_\_.mpls.0**. When the egress node fails, the PLR or the Protector will switch traffic

to this backup next hop in PFE. The outer label (the transport LSP label) of packets is popped, and the inner label (the layer 3 VPN label allocated by the egress node) is looked up in `__context__mpls.0`, which results in forwarding the packets directly to the CE (in Collocated protector model) or the backup PE (in Centralized protector model).

For more information about egress PE failure protection, see Internet draft [draft-minto-2547-egress-node-fast-protection-00](#), *2547 egress PE Fast Failure Protection*.

### IGP Advertisement Model

---

Egress protection availability is advertised in the interior gateway protocol (IGP). Label protocols along with Constrained Shortest Path First (CSPF) use this information to do egress protection.

For Layer 3 VPNs, the IGP advertisements can be of the following types:

- Context identifier as a stub link (supported in Junos OS 11.4 R3 and later). A link connecting a stub node to a transit node is a stub link.
- Context identifier as a stub alias node (supported in Junos OS 13.3 and later).
- Context identifier as a stub proxy node (supported in Junos OS 13.3 and later).

By default, the stub link is used. To enable enhanced point-of-local-repair (PLR) functionality, in which the PLR reroutes service traffic during an egress failure, configure a stub alias node or a stub proxy node as follows:

```
[edit protocols mpls egress-protection context-identifier 192.0.2.6]
user@host# set advertise-mode ?
Possible completions:
  stub-alias      Alias
  stub-proxy      Proxy
```

The two methods offer different advantages, depending on the needs of your network deployment.

#### **Context Identifier as a Stub Alias Node**

In the stub alias method, the LSP end-point address has an explicit backup egress node where the backup can be learned or configured on the penultimate hop node of a protected LSP. With this model, the penultimate hop node of a protected LSP sets up the bypass LSP tunnel to back up the egress node by avoiding the primary egress node. This model requires a Junos OS upgrade in core nodes, but is flexible enough to support all traffic engineering constraints.

The PLR learns that the context ID has a protector. When the primary context ID goes down, packets are rerouted to the protector by way of a pre-programmed backup path. The context ID and protector mapping are configured or learned on the PLR and signaled in the IGP from the protector. A routing table called `inet.5` on the PLR provides the configured or IGP-learned details.

IS-IS advertises context IDs into the TED through an IP address TLV. IS-IS imports this TLV into the TED as extended information. IS-IS advertises the protector TLV routes in the `inet.5` route for the context ID with protocol next hop being the protector's router ID.

If the protector TLV has a label, the label is added to the route in the inet.5 routing table for LDP to use.

CSPF considers the IP address TLV for tunnel endpoint computation.

With the stub alias model, the protector LSP setup does not require any changes in any nodes. But bypass LSP setup for node protection requires changes in the PHN and the protector router.

When RSVP sets up bypass for node protection LSP, RSVP also performs a lookup for the protector if the PLR is the penultimate hop of the LSP. If the protector is available for the LSP destination, it uses CSPF to compute a path with a constraint that excludes the egress PE and sets up a bypass LSP destination to the context ID if one is not already set up. When setting up a bypass LSP to the context ID, the PLR unsets all protection options.

LDP is useful in the case when the network supports 100 percent LFA coverage but does not support 100 percent per-prefix LFA coverage. LDP sets up a backup path with the protector with the context label advertised by the protector to the service point.

In networks in which 100 percent LFA coverage is not available, it is useful to have backup LSP LFAs with RSVP-based tunnels.

In a steady state, the forwarding is the same as on any other protected LSP in the PLR. In the protector, the non-null label that is advertised and signaled for the context ID has the table next hop point to the MPLS context table, where the peers' labels are programmed.

During a failure, the PLR swaps the transport label with the bypass LSP for the context ID or swaps the label context-label (the protector-advertised label for the context ID) and pushes the transport label to the protector lo0 interface address.

### ***Context Identifier as a Stub Proxy Node***

Context identifier as a stub proxy node (supported in Junos OS 13.3 and later). A stub node is one that only appears at the end of an AS path, which means it does not provide transit service. In this mode, known as the virtual or proxy mode, the LSP end-point address is represented as a node with bidirectional links, with the LSP's primary egress node and backup egress node. With this representation, the penultimate hop of the LSP primary egress point can behave like a PLR in setting up a bypass tunnel to back up the egress by avoiding the primary egress node. This model has the advantage that you do not need to upgrade Junos OS on core nodes and will thereby help operators to deploy this technology.

The context ID is represented as a node in the traffic engineering (TE) and IGP databases. The primary PE device advertises the context node into the IGP and TE databases. The primary PE device and the protected PE device support one link to the context node with a bandwidth and a TE metric. Other TE characteristics of TE links are not advertised by Junos OS.

In IS-IS, the primary PE router advertises the proxy node along with links to the primary router and the protector router. The primary and the protector routers advertise links to the proxy node. The proxy node builds the following information.

- System ID—Binary-coded decimal based on the context ID.
- Host name—Protector-name:context ID
- LSP-ID—<System-ID>.00
- PDU type—Level 2 and Level 1, based on the configuration
- LSP attributes:
  - Overload—1
  - IS\_TYPE\_L1(0x01) | IS\_TYPE\_L2(0x02) for the level 2 PDU
  - IS\_TYPE\_L1 for level 1
  - Multiarea—No
  - All other attributes—0

The proxy node only contains area, MT, host name, router ID, protocols and IS reachability TLVs. The area, MT, authentication, and protocols TLV are the same as on the primary. The IS reachability TLVs contains two links called Cnode-primary-link and Cnode-protector-link. Both links include TE TLVs. The following TE-link-TLVs are advertised in context links:

- IPv4 interface or neighbor address
- Maximum bandwidth
- TE default metric
- Link (local or remote) Identifiers

Sub TLV values:

- Bandwidth—zero
- TE metric—Maximum TE metric
- Interface address—context ID
- Protector neighbor address—protector router ID
- Primary neighbor address—protected router ID
- Link local-ID protector—0x80fffff1
- Link local-ID primary—0x80fffff2
- Link remote-ID protector—Learned from protector
- Link remote-ID primary—Learned from primary

Protected PE links to context node (primary advertises the link with the following details):

- Bandwidth—Maximum
- TE metric—1
- Interface address—Router ID
- Context neighbor address—Context ID
- Link local-ID to context node—Automatically generated (similar to a sham link)
- Link remote-ID to context node—0x80ffff2

Protector PE links to context node:

- The protector advertises unnumbered transit links with the maximum routable link metric and the maximum TE metric and zero bandwidth to the context node. Other TE characteristics are not advertised.

Unnumbered links are advertised with the following attributes:

- bandwidth—0
- TE metric—MAX TE metric
- Interface address—Router ID
- Context neighbor address—Context ID
- Link local ID to context node—Autogenerated (similar to a sham link)
- Link remote ID to context node—0x80ffff1

In RSVP, the behavior changes are only in the protector and primary routers. RSVP terminates the LSP and the bypass LSP to the context ID. If the context ID is the protector, a non-null label is signaled. Otherwise, it will be based on the configuration or the requested label type. RSVP verifies the Explicit Route Object (ERO) from the path for itself and the context ID. RSVP sends the Resv message with two Record Route Object (RRO) objects—one for the context ID and one for itself. This simulates the penultimate-hop node (PHN) to do node protection with the protector for the primary for context ID LSP. As the fast reroute (FRR)-required bypass, the LSP has to merge back to the protector LSP PHN setup bypass to context ID through the protector by avoiding the primary.

The protector also terminates the backup LSP for the context ID to keep the protected LSP alive during a failure until the ingress node resigns the LSP. The new LSP is reestablished through the protector, but this LSP is not used for service traffic as service protocol does not use the context ID. The LSP traverses through the protector even if the primary comes up. Only reoptimization resigns the LSP through the primary. In stub proxy mode, the bypass LSP with constraints is not supported.

LDP cannot use the stub proxy method due to the inflated metric advertised in the IGP.

With regard the forwarding state, a PE router that protects one or more segments that are connected to another PE is referred to as a protector PE. A protector PE must learn the forwarding state of the segments that it is protecting from the primary PE that is being protected.

For a given segment, if the protector PE is not directly connected to the CE device associated with the segment, it must also learn the forwarding state from at least one backup PE. This situation might arise only in the case of egress PE failure protection.

A protector PE maintains forwarding state for a given segment in the context of the primary PE. A protector PE might maintain state for only a subset of the segments on the primary PE or for all the segments on the primary PE.

- See Also**
- [Example: Configuring Egress Protection for Layer 3 VPN Services on page 834](#)
  - [Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP on page 843](#)

## Example: Configuring Egress Protection for Layer 3 VPN Services

This example shows how to configure egress protection for fast restoration of Layer 3 VPN services.

- [Requirements on page 834](#)
- [Overview on page 834](#)
- [Configuration on page 836](#)
- [Verification on page 841](#)

### Requirements

---

This example uses the following hardware and software components

- MX Series 3D Universal Edge Routers
- Tunnel PICs or the configuration of the Enhanced IP Network Services mode (using the **network-services enhanced-ip** statement at the **[edit chassis]** hierarchy level).
- Junos OS Release 11.4R3 or later running on the devices

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Configuration Guide*.
- Configure the following routing protocols on all the PE and PLR routers.
  - MPLS, LSPs, and LDP. See the *Junos OS MPLS Applications Configuration Guide*.
  - BGP and IS-IS. See the *Junos OS Routing Protocols Configuration Guide*.
- Configure Layer 3 VPNs. See the *Junos OS VPNs Configuration Guide*.

### Overview

---

Typically, Layer 3 VPN service restoration, in case of egress PE router failure (for multihomed customer edge [CE] routers), depends on the ingress PE router to detect the egress PE node failure and switch traffic to the backup PE router for multihomed CE sites.

Junos OS Release 11.4R3 or later enables you to configure egress protection for Layer 3 VPN services that protects the services from egress PE node failure in a scenario where the CE site is multihomed with more than one PE router. The mechanism enables local repair to be performed immediately upon an egress node failure. The router acting as the point of local repair (PLR) redirects VPN traffic to a protector PE router for restoring service quickly, achieving fast protection that is comparable to MPLS fast reroute.

The statements used to configure egress protection are:

- **egress-protection**—When configured at the `[edit protocols mpls]` hierarchy level, this statement specifies protector information and the context identifier for the Layer 3 VPN and edge protection virtual circuit:

```
[edit protocols mpls]
egress-protection {
  context-identifier context-id {
    primary | protector;
    metric igp-metric-value;
  }
}
```

When configured at the `[edit protocols bgp group group-name family inet-vpn unicast]`, `[edit protocols bgp group group-name family inet6-vpn unicast]`, or `[edit protocols bgp group group-name family iso-vpn unicast]` hierarchy levels, the `egress-protection` statement specifies the context identifier that enables egress protection for the configured BGP VPN network layer reachability information (NLRI).

```
[edit protocols bgp]
group internal {
  type internal;
  local-address ip-address;
  family <inet-vpn|inet6-vpn|iso-vpn> {
    unicast {
      egress-protection {
        context-identifier {
          context-id-ip-address;
        }
      }
    }
  }
}
```

When configured at the `[edit routing-instances]` hierarchy level, the `egress-protection` statement holds the context identifier of the protected PE router.

This configuration must be done only in the primary PE router and is used for outbound BGP updates for the next hops.

```
[edit routing-instance]
routing-instance-name {
  egress-protection {
    context-identifier {
      context-id-ip-address;
    }
  }
}
```

Configuring the **context-identifier** statement at the **[edit routing-instances routing-instance-name]** hierarchy level provides customer edge VRL-level context ID granularity for each VRF instance.

- **context-identifier**—This statement specifies an IPv4 address used to define the pair of PE routers participating in the egress protection LSP. The context identifier is used to assign an identifier to the protector PE router. The identifier is propagated to the other PE routers participating in the network, making it possible for the protected egress PE router to signal the egress protection LSP to the protector PE router.

## Configuration

### CLI Quick Configuration



**NOTE:** This example only shows sample configuration that is relevant to configuring egress PE protection for Layer 3 VPN services on the protected router, PE2, the protector router, PE3, and the PLR router.

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.

PE2 (Protected PE Router)	<pre> set protocols mpls interface all set protocols mpls interface fxp0.0 disable set protocols mpls egress-protection context-identifier 192.0.2.6 primary set protocols bgp group ibgp type internal set protocols bgp group ibgp local-address 10.255.245.194 set protocols bgp group ibgp family inet-vpn unicast egress-protection context-identifier 192.0.2.6 </pre>
PE3 (Protector PE Router)	<pre> set protocols mpls interface all set protocols mpls interface fxp0.0 disable set protocols mpls egress-protection context-identifier 192.0.2.6 protector set protocols bgp group ibgp type internal set protocols bgp group ibgp local-address 10.255.245.196 set protocols bgp group ibgp family inet-vpn unicast egress-protection keep-import remote-vrf set policy-options policy-statement remote-vrf from community rsite1 set policy-options policy-statement remote-vrf from community rsite24 set policy-options policy-statement remote-vrf then accept set policy-options community rsite1 members target:1:1 set policy-options community rsite24 members target:100:1023 </pre>
PLR Router	<pre> set protocols mpls interface all set protocols mpls interface fxp0.0 disable set protocols isis level 1 disable set protocols isis interface all node-link-protection set protocols isis backup-spf-options per-prefix-calculation set protocols ldp track-igp-metric set protocols ldp interface all </pre>



```
set protocols ldp interface fxp0.0 disable
```

### *Configuring the Protected PE Router (PE2)*

**Step-by-Step Procedure** To configure the protected PE router, PE2:

1. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
```

2. Configure egress protection and the context identifier.



**NOTE:** The context identifier type must be set to primary.

```
[edit protocols mpls]
user@PE2# set egress-protection context-identifier 192.0.2.6 primary
```

3. Configure egress protection for the configured BGP NRLI.



**NOTE:** The context identifier configured at the [edit protocols bgp group group-name family inet-vpn] hierarchy level should match the context identifier configured at the [edit protocols mpls] hierarchy level.

```
[edit protocols bgp]
user@PE2# set group ibgp type internal
user@PE2# set group ibgp local-address 10.255.245.194
user@PE2# set group ibgp family inet-vpn unicast egress-protection
context-identifier 192.0.2.6
```



**NOTE:** Configuring the context-identifier at the [edit routing-instances routing-instance-name] hierarchy level provides CE VRF-level context-id granularity for each virtual routing and forwarding (VRF) instance.

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

```
[edit]
user@PE2# commit
```

**Results** Confirm your configuration by issuing the show protocols command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE2# show protocols
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 192.0.2.6 {
      primary;
    }
  }
}
bgp {
  group ibgp {
    type internal;
    local-address 10.255.245.194;
    family inet-vpn {
      unicast {
        egress-protection {
          context-identifier {
            192.0.2.6;
          }
        }
      }
    }
  }
}
}

```

### *Configuring the Protector PE Router (PE3)*

#### **Step-by-Step Procedure**

To configure the protector PE router, PE3:

1. Configure MPLS on the interfaces.  

```

[edit protocols mpls]
user@PE3# set interface all
user@PE3# set mpls interface fxp0.0 disable

```
2. Configure egress protection and the context identifier.  

```

[edit protocols mpls]
user@PE3# set egress-protection context-identifier 192.0.2.6 protector

```
3. Configure IPv4 Layer 3 VPN NRLI parameters.  

```

[edit protocols bgp]
user@PE3# set group ibgp type internal
user@PE3# set group ibgp local-address 10.255.245.196
user@PE3# set group ibgp family inet-vpn unicast egress-protection keep-import remote-vrf

```
4. Configure routing policy options.  

```

[edit policy-options]
user@PE3# set policy-statement remote-vrf from community rsite1

```

```

user@PE3# set policy-statement remote-vrf from community rsite24
user@PE3# set policy-statement remote-vrf then accept
user@PE3# set community rsite1 members target:1:1
user@PE3# set community rsite24 members target:100:1023

```

5. After you are done configuring the device, commit the configuration.

```

[edit]
user@PE3# commit

```

**Results** Confirm your configuration by issuing the **show protocols** and the **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE3# show protocols
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 192.0.2.6 {
      protector;
    }
  }
}
bgp {
  group ibgp {
    type internal;
    local-address 10.255.245.196;
    family inet-vpn {
      unicast {
        egress-protection {
          keep-import remote-vrf;
        }
      }
    }
  }
}
}

user@PE3# show policy-options
policy-statement remote-vrf {
  from community [ rsite1 rsite24 ];
  then accept;
}
community rsite1 members target:1:1;
community rsite24 members target:100:1023;

```

### *Configuring the PLR Router*

**Step-by-Step Procedure** To configure the router acting as the point of local repair (PLR):

1. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PLR# set interface all
user@PLR# set interface fxp0.0 disable
```

2. Configure per-prefix-LFA calculation along with link protection.

```
[edit protocols isis]
user@PLR# set backup-spf-options per-prefix-calculation
user@PLR# set level 1 disable
user@PLR# set interface all node-link-protection
user@PLR# set interface fxp0.0 disable
```

3. Configure LDP to use the interior gateway protocol (IGP) route metric instead of the default LDP route metric (the default LDP route metric is 1).

```
[edit protocols ldp]
user@PLR# set track-igp-metric
user@PLR# set interface all
user@PLR# set interface fxp0.0 disable
```

**Results** Confirm your configuration by issuing the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PLR# show protocols
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
isis {
  backup-spf-options per-prefix-calculation;
  level 1 disable;
  interface all {
    node-link-protection;
  }
}
ldp {
  track-igp-metric;
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying Egress Protection Details on page 841](#)
- [Verifying Routing Instances on page 841](#)
- [Verifying BGP NRLI on page 842](#)

### *Verifying Egress Protection Details*

**Purpose** Check the egress protection configuration.

**Action** user@PE3> show mpls egress-protection details

```
Instance           Type           Protection-Type
rsite1             remote-vrf    Protector
  RIB __192.0.2.6-rsite1__.inet.0, Context-Id 192.0.2.6, Enhanced-lookup
  Route Target 1:1
rsite24            remote-vrf    Protector
  RIB __192.0.2.6-rsite24__.inet.0, Context-Id 192.0.2.6, Enhanced-lookup
  Route Target 100:1023
```

**Meaning** **Instance** indicates the routing-instance name. **Type** shows the type of the VRF. It can be either **local-vrf** or **remote-vrf**. **RIB** (routing information base) indicates the edge-protection created routing table. **Context-Id** shows the context ID associated with the RIB. **Route Target** shows the route target associated with the routing instance.

### *Verifying Routing Instances*

**Purpose** Verify the routing instances.

**Action** user@PE3> show route instance site1 detail

```
site1:
  Router ID: 198.51.100.1
  Type: vrf                               State: Active
  Interfaces:
    lt-1/3/0.8
  Route-distinguisher: 10.255.255.11:150
  Vrf-import: [ site1-import ]
  Vrf-export: [ __vrf-export-site1-internal__ ]
  Vrf-export-target: [ target:100:250 ]
  Fast-reroute-priority: low
  Vrf-edge-protection-id: 192.0.2.6
  Tables:
    site1.inet.0          : 27 routes (26 active, 0 holddown, 0 hidden)
    site1.iso.0           : 0 routes (0 active, 0 holddown, 0 hidden)
    site1.inet6.0         : 0 routes (0 active, 0 holddown, 0 hidden)
    site1.mdt.0           : 0 routes (0 active, 0 holddown, 0 hidden)
```

**Meaning** **Vrf-edge-protection-id** shows the egress protection configured in the protector PE router with the routing instance.

#### *Verifying BGP NLRI*

**Purpose** Check the details of the BGP VPN network layer reachability information.

**Action** user@PE3> show bgp neighbor

```
Peer: 10.255.55.1+179 AS 65535 Local: 10.255.22.1+59264 AS 65535
  Type: Internal   State: Established   Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress KeepAll AddressFamily Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 10.255.22.1 Holdtime: 90 Preference: 170
  NLRI configured with egress-protection: inet-vpn-unicast
  Egress-protection NLRI inet-vpn-unicast, keep-import: [ VPN-A-remote ]
  Number of flaps: 0
```

**Meaning** **NLRI configured with egress-protection** shows the BGP family configured with egress protection. **egress-protection NLRI inet-vpn-unicast, keep-import: [remote-vrf]** shows the egress protection routing policy for the BGP group.

**See Also** • [Egress Protection for Layer 3 VPN Edge Protection Overview on page 827](#)

## Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

This example shows how to configure fast service restoration at the egress of a Layer 3 VPN when the customer is multihomed to the service provider. Further, this example includes enhanced point-of-local-repair (PLR) functionality, in which the PLR reroutes service traffic during an egress failure.

Starting in Junos OS Release 13.3, enhanced PLR functionality is available, in which the PLR reroutes service traffic during an egress failure. As part of this enhancement, the PLR router no longer needs to be directly connected to the protector router. Previously, if the PLR was not directly connected to the protector router, the loop-free alternate route could not find the backup path to the protector.

- [Requirements on page 843](#)
- [Overview on page 843](#)
- [Configuration on page 844](#)
- [Verification on page 859](#)

### Requirements

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

This example requires Junos OS Release 13.3 or later.

### Overview

In this example, the customer edge (CE) devices are part of a VPN where Device CE1 is multihomed with Device PE2 and Device PE3.

Device PE3 acts as the protector for the Layer 3 VPN routing instances or subnets.

Device PE1 is the originator for the context identifier for Device CE1, Device PE2 is the primary router for that context identifier, while Device PE3 is the protector for that context identifier.

Device P1 acts as the point of local repair (PLR). As such, Device P1 can redirect Layer 3 VPN traffic to the protector PE router to enable fast restoration and reroute.

The working path is through P1>PE2. The backup path is through P1>PE3. Traffic flows through the working path under normal circumstances. When a Device PE2 node or link failure is detected, traffic is rerouted from the working path to the protected path. In the normal failover process, the detection of failure and the recovery rely on the control plane and is therefore relatively slow. Typically, if there is a link or node failure in the core network, the egress PE router would have to rely on the ingress PE router to detect the failure and switch over to the backup path, because a local repair option for egress failure is not available. To provide a local repair solution for the egress PE link or node failure, a mechanism known as egress protection is used in this example to repair and restore the connection quickly. Because egress protection is configured, the PLR router detects the Device PE2 link or node failure and reroutes traffic through the protector Device PE3 using


the backup LDP-signaled label-switched path (LSP). The PLR router uses per-prefix loop-free alternate routes to program the backup next hop through Device PE3, and traffic is forwarded to Device CE2 using the alternate paths. This restoration is done quickly after the PLR router detects the Device PE2 egress node or link failure. The dual protection mechanism can also be used for egress protection where the two PE routers can simultaneously act as the primary PE router and the protector PE router for their respective context ID routes or next hops.

In addition to egress protection, this example demonstrates an enhanced PLR function, in which the PLR reroutes service traffic during the egress failure. This enhancement is supported in Junos OS Release 13.3 and later. In this example, Device P1 (the PLR) is directly connected to Device PE3 (the protector). A new configuration statement, **advertise-mode**, enables you to set the method for the interior gateway protocol (IGP) to advertise egress protection availability.

## Topology

Figure 81 on page 844 shows the sample network.

Figure 81: Layer 3 VPN Egress Protection with RSVP and LDP



## Configuration

<b>CLI Quick Configuration</b>	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 <b>[edit]</b> hierarchy level.
<b>Device CE1</b>	<pre> set interfaces ge-1/2/0 unit 0 description to_PE1 set interfaces ge-1/2/0 unit 0 family inet address 10.1.0.1/24 set interfaces lo0 unit 0 family inet address 172.16.0.1/32 set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 </pre>
<b>Device CE2</b>	<pre> set interfaces ge-1/2/2 unit 0 description to_PE2 set interfaces ge-1/2/2 unit 0 family inet address 10.8.0.2/24 set interfaces ge-1/2/0 unit 0 description to_PE3 set interfaces ge-1/2/0 unit 0 family inet address 10.9.0.2/24 set interfaces lo0 unit 0 family inet address 172.16.0.2/32 set protocols ospf area 0.0.0.0 interface ge-1/2/2.0 set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 </pre>



```

Device P1      set interfaces ge-1/2/1 unit 0 description to_PE1
                set interfaces ge-1/2/1 unit 0 family inet address 10.2.0.2/24
                set interfaces ge-1/2/1 unit 0 family iso
                set interfaces ge-1/2/1 unit 0 family mpls
                set interfaces ge-1/2/0 unit 0 description to_PE2
                set interfaces ge-1/2/0 unit 0 family inet address 10.4.0.1/24
                set interfaces ge-1/2/0 unit 0 family iso
                set interfaces ge-1/2/0 unit 0 family mpls
                set interfaces ge-1/2/2 unit 0 description to_PE3
                set interfaces ge-1/2/2 unit 0 family inet address 10.5.0.1/24
                set interfaces ge-1/2/2 unit 0 family iso
                set interfaces ge-1/2/2 unit 0 family mpls
                set interfaces lo0 unit 0 family inet address 172.16.0.3/32
                set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0003.00
                set protocols rsvp interface all
                set protocols rsvp interface fxp0.0 disable
                set protocols mpls interface all
                set protocols isis backup-spf-options per-prefix-calculation
                set protocols isis level 1 disable
                set protocols isis level 2 wide-metrics-only
                set protocols isis interface all node-link-protection
                set protocols isis interface fxp0.0 disable
                set protocols isis interface lo0.0
                set protocols ldp track-igp-metric
                set protocols ldp interface all
                set protocols ldp interface fxp0.0 disable

Device PE1     set interfaces ge-1/2/0 unit 0 description to_CE1
                set interfaces ge-1/2/0 unit 0 family inet address 10.1.0.2/24
                set interfaces ge-1/2/1 unit 0 description to_P1
                set interfaces ge-1/2/1 unit 0 family inet address 10.2.0.1/24
                set interfaces ge-1/2/1 unit 0 family iso
                set interfaces ge-1/2/1 unit 0 family mpls
                set interfaces lo0 unit 0 family inet address 172.16.183.55/32
                set interfaces lo0 unit 0 family iso address 49.0002.1720.1618.3055.00
                set protocols rsvp interface all
                set protocols rsvp interface fxp0.0 disable
                set protocols mpls label-switched-path toPrimary192.0.2.6 to 192.0.2.6
                set protocols mpls label-switched-path toPrimary192.0.2.6 egress-protection
                set protocols mpls interface all
                set protocols bgp group ibgp type internal
                set protocols bgp group ibgp local-address 172.16.183.55
                set protocols bgp group ibgp family inet-vpn unicast
                set protocols bgp group ibgp neighbor 172.16.183.56
                set protocols bgp group ibgp neighbor 172.16.183.59
                set protocols isis level 1 disable
                set protocols isis level 2 wide-metrics-only
                set protocols isis interface all
                set protocols isis interface fxp0.0 disable
                set protocols isis interface lo0.0
                set protocols ldp track-igp-metric
                set protocols ldp interface all
                set protocols ldp interface fxp0.0 disable
                set routing-instances vpn1 instance-type vrf
                set routing-instances vpn1 interface ge-1/2/0.0

```

```

set routing-instances vpn1 route-distinguisher 172.16.183.55:10
set routing-instances vpn1 vrf-target target:10:10
set routing-instances vpn1 routing-options static route 100.0.0.0/24 next-hop 10.1.0.1
set routing-instances vpn1 protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set routing-options autonomous-system 64510

```

Device PE2

```

set interfaces ge-1/2/0 unit 0 description to_P1
set interfaces ge-1/2/0 unit 0 family inet address 10.4.0.2/24
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 description to_CE2
set interfaces ge-1/2/2 unit 0 family inet address 10.8.0.1/24
set interfaces ge-1/2/1 unit 0 description to_PE3
set interfaces ge-1/2/1 unit 0 family inet address 10.7.0.1/24
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 172.16.183.56/32
set interfaces lo0 unit 0 family iso address 49.0002.1720.1618.3056.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path toPE1 to 172.16.183.55
set protocols mpls label-switched-path toPrimary192.0.2.6 to 192.0.2.6
set protocols mpls label-switched-path toPrimary192.0.2.6 egress-protection
set protocols mpls interface all
set protocols mpls egress-protection context-identifier 192.0.2.6 primary
set protocols mpls egress-protection context-identifier 192.0.2.6 advertise-mode
  stub-proxy
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.183.56
set protocols bgp group ibgp family inet-vpn unicast egress-protection context-identifier
  192.0.2.6
set protocols bgp group ibgp neighbor 172.16.183.55
set protocols bgp group ibgp neighbor 172.16.183.59
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0
set protocols ldp track-igp-metric
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-options autonomous-system 64510

```

Device PE3

```

set interfaces ge-1/2/2 unit 0 description to_P1
set interfaces ge-1/2/2 unit 0 family inet address 10.5.0.2/24
set interfaces ge-1/2/2 unit 0 family iso
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces ge-1/2/0 unit 0 description to_CE2
set interfaces ge-1/2/0 unit 0 family inet address 10.9.0.1/24
set interfaces ge-1/2/1 unit 0 description to_PE2
set interfaces ge-1/2/1 unit 0 family inet address 10.7.0.2/24
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 172.16.183.59/32

```

```

set interfaces lo0 unit 0 family iso address 49.0002.1720.1618.3059.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path toPE1 to 172.16.183.55
set protocols mpls interface all
set protocols mpls egress-protection context-identifier 192.0.2.6 protector
set protocols mpls egress-protection context-identifier 192.0.2.6 advertise-mode
  stub-proxy
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.183.59
set protocols bgp group ibgp family inet-vpn unicast egress-protection keep-import
  remote-vrf
set protocols bgp group ibgp neighbor 172.16.183.55
set protocols bgp group ibgp neighbor 172.16.183.56
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0
set protocols ldp track-igp-metric
set protocols ldp interface all
set policy-options policy-statement remote-vrf from community rsite1
set policy-options policy-statement remote-vrf from community rsite24
set policy-options policy-statement remote-vrf then accept
set policy-options community rsite1 members target:1:1
set policy-options community rsite24 members target:100:1023
set routing-options autonomous-system 64510

```

**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 P1 (the PLR):

1. Configure the device interfaces.

```

[edit interfaces]
user@P1# set ge-1/2/1 unit 0 description to_PE1
user@P1# set ge-1/2/1 unit 0 family inet address 10.2.0.2/24
user@P1# set ge-1/2/1 unit 0 family iso
user@P1# set ge-1/2/1 unit 0 family mpls

```

```

user@P1# set ge-1/2/0 unit 0 description to_PE2
user@P1# set ge-1/2/0 unit 0 family inet address 10.4.0.1/24
user@P1# set ge-1/2/0 unit 0 family iso
user@P1# set ge-1/2/0 unit 0 family mpls

```

```

user@P1# set ge-1/2/2 unit 0 description to_PE3
user@P1# set ge-1/2/2 unit 0 family inet address 10.5.0.1/24
user@P1# set ge-1/2/2 unit 0 family iso
user@P1# set ge-1/2/2 unit 0 family mpls

```

```

user@P1# set lo0 unit 0 family inet address 172.16.0.3/32
user@P1# set lo0 unit 0 family iso address 49.0002.0172.0016.0003.00

```

## 2. Configure IS-IS.

Configure per-prefix-LFA calculation along with node link protection.

```
[edit protocols isis]
user@P1# set backup-spf-options per-prefix-calculation
user@P1# set level 1 disable
user@P1# set level 2 wide-metrics-only
user@P1# set interface all node-link-protection
user@P1# set interface fxp0.0 disable
user@P1# set interface lo0.0
```

## 3. Enable MPLS.

```
[edit protocols mpls ]
user@P1# set interface all
```

## 4. Enable RSVP.

```
[edit protocols rsvp]
user@P1# set interface all
user@P1# set interface fxp0.0 disable
```

## 5. Enable LDP.

```
[edit protocols ldp]
user@P1# set track-igp-metric
user@P1# set interface all
user@P1# set interface fxp0.0 disable
```

**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 PE1:

## 1. Configure the device interfaces.

```
[edit interfaces]
user@PE1# set ge-1/2/0 unit 0 description to_CE1
user@PE1# set ge-1/2/0 unit 0 family inet address 10.1.0.2/24

user@PE1# set ge-1/2/1 unit 0 description to_P1
user@PE1# set ge-1/2/1 unit 0 family inet address 10.2.0.1/24
user@PE1# set ge-1/2/1 unit 0 family iso
user@PE1# set ge-1/2/1 unit 0 family mpls

user@PE1# set lo0 unit 0 family inet address 172.16.183.55/32
user@PE1# set lo0 unit 0 family iso address 49.0002.1720.1618.3055.00
```

## 2. Enable RSVP.

```
[edit protocols rsvp]
```

```
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
```

3. Configure MPLS.

```
[edit protocols mpls]
user@PE1# set label-switched-path toPrimary192.0.2.6 to 192.0.2.6
user@PE1# set label-switched-path toPrimary192.0.2.6 egress-protection
user@PE1# set interface all
```

4. Configure IBGP.

```
[edit protocols bgp group ibgp]
user@PE1# set type internal
user@PE1# set local-address 172.16.183.55
user@PE1# set family inet-vpn unicast
user@PE1# set neighbor 172.16.183.56
user@PE1# set neighbor 172.16.183.59
```

5. Configure IS-IS.

```
[edit protocols isis]
user@PE1# set level 1 disable
user@PE1# set level 2 wide-metrics-only
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
user@PE1# set interface lo0.0
```

6. Enable LDP.

```
[edit protocols ldp]
user@PE1# set track-igp-metric
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
```

7. Configure the routing instance.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-1/2/0.0
user@PE1# set route-distinguisher 172.16.183.55:10
user@PE1# set vrf-target target:10:10
user@PE1# set routing-options static route 100.0.0.0/24 next-hop 10.1.0.1
user@PE1# set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
```

8. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set autonomous-system 64510
```

**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 PE2:

1. Configure the device interfaces.

```
[edit interfaces]
user@PE2# set ge-1/2/0 unit 0 description to_P1
user@PE2# set ge-1/2/0 unit 0 family inet address 10.4.0.2/24
user@PE2# set ge-1/2/0 unit 0 family iso
user@PE2# set ge-1/2/0 unit 0 family mpls

user@PE2# set ge-1/2/2 unit 0 description to_CE2
user@PE2# set ge-1/2/2 unit 0 family inet address 10.8.0.1/24

user@PE2# set ge-1/2/1 unit 0 description to_PE3
user@PE2# set ge-1/2/1 unit 0 family inet address 10.7.0.1/24
user@PE2# set ge-1/2/1 unit 0 family iso
user@PE2# set ge-1/2/1 unit 0 family mpls

user@PE2# set lo0 unit 0 family inet address 172.16.183.56/32
user@PE2# set lo0 unit 0 family iso address 49.0002.1720.1618.3056.00
```

2. Enable RSVP.

```
[edit protocols rsvp]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
```

3. Configure MPLS.

```
[edit protocols mpls]
user@PE2# set label-switched-path toPE1 to 172.16.183.55
user@PE2# set label-switched-path toPrimary192.0.2.6 to 192.0.2.6
user@PE2# set label-switched-path toPrimary192.0.2.6 egress-protection
user@PE2# set interface all
user@PE2# set egress-protection context-identifier 192.0.2.6 primary
user@PE2# set egress-protection context-identifier 192.0.2.6 advertise-mode
stub-proxy
```

4. Configure IBGP.

```
[edit protocols bgp group ibgp]
user@PE2# set type internal
user@PE2# set local-address 172.16.183.56
user@PE2# set family inet-vpn unicast egress-protection context-identifier 192.0.2.6
user@PE2# set neighbor 172.16.183.55
user@PE2# set neighbor 172.16.183.59
```

5. Configure IS-IS.

```
[edit protocols isis]
```

```

user@PE2# set level 1 disable
user@PE2# set level 2 wide-metrics-only
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
user@PE2# set interface lo0.0

```

6. Enable LDP.

```

[edit protocols ldp]
user@PE2# set track-igp-metric
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable

```

7. Configure the AS number.

```

[edit routing-options]
user@PE2# set autonomous-system 64510

```

**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 PE3:

1. Configure the device interfaces.

```

[edit interfaces]
user@PE3# set ge-1/2/2 unit 0 description to_P1
user@PE3# set ge-1/2/2 unit 0 family inet address 10.5.0.2/24
user@PE3# set ge-1/2/2 unit 0 family iso
user@PE3# set ge-1/2/2 unit 0 family mpls

user@PE3# set ge-1/2/0 unit 0 description to_CE2
user@PE3# set ge-1/2/0 unit 0 family inet address 10.9.0.1/24

user@PE3# set ge-1/2/1 unit 0 description to_PE2
user@PE3# set ge-1/2/1 unit 0 family inet address 10.7.0.2/24
user@PE3# set ge-1/2/1 unit 0 family iso
user@PE3# set ge-1/2/1 unit 0 family mpls

user@PE3# set lo0 unit 0 family inet address 172.16.183.59/32
user@PE3# set lo0 unit 0 family iso address 49.0002.1720.1618.3059.00

```

2. Enable RSVP.

```

[edit protocols rsvp]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable

```

3. Configure MPLS.

```

[edit protocols mpls]

```

```

user@PE3# set label-switched-path toPE1 to 172.16.183.55
user@PE3# set interface all
user@PE3# set egress-protection context-identifier 192.0.2.6 protector
user@PE3# set egress-protection context-identifier 192.0.2.6 advertise-mode
stub-proxy

```

4. Configure IBGP.

```

[edit protocols bgp group ibgp]
user@PE3# set type internal
user@PE3# set local-address 172.16.183.59
user@PE3# set family inet-vpn unicast egress-protection keep-import remote-vrf
user@PE3# set neighbor 172.16.183.55
user@PE3# set neighbor 172.16.183.56

```

5. Configure IS-IS.

```

[edit protocols isis]
user@PE3# set level 1 disable
user@PE3# set level 2 wide-metrics-only
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
user@PE3# set interface lo0.0

```

6. Enable LDP.

```

[edit protocols ldp]
user@PE3# set track-igp-metric
user@PE3# set interface all

```

7. Configure the routing policy.

```

[edit policy-options]
user@PE3# set policy-statement remote-vrf from community rsite1
user@PE3# set policy-statement remote-vrf from community rsite24
user@PE3# set policy-statement remote-vrf then accept
user@PE3# set community rsite1 members target:1:1
user@PE3# set community rsite24 members target:100:1023

```

8. Configure the AS number.

```

[edit routing-options]
user@PE3# set autonomous-system 64510

```

**Results** From configuration mode, 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.

**Device P1**

```

user@P1# show interfaces
ge-1/2/0 {
  unit 0 {
    description to_PE2;

```



```

        family inet {
            address 10.4.0.1/24;
        }
        family iso;
        family mpls;
    }
}
ge-1/2/1 {
    unit 0 {
        description to_PE1;
        family inet {
            address 10.2.0.2/24;
        }
        family iso;
        family mpls;
    }
}
ge-1/2/2 {
    unit 0 {
        description to_PE3;
        family inet {
            address 10.5.0.1/24;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.0.3/32;
        }
        family iso {
            address 49.0002.0172.0016.0003.00;
        }
    }
}
}

```

user@P1# show protocols

```

rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
}
isis {
    backup-spf-options per-prefix-calculation;
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        node-link-protection;
    }
    interface fxp0.0 {

```

```

        disable;
    }
    interface lo0.0;
}
ldp {
    track-igp-metric;
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```

**Device PE1**

```

user@PE1# show interfaces
ge-1/2/0 {
    unit 0 {
        description to_CE1;
        family inet {
            address 10.1.0.2/24;
        }
    }
}
ge-1/2/1 {
    unit 0 {
        description to_P1;
        family inet {
            address 10.2.0.1/24;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.183.55/32;
        }
        family iso {
            address 49.0002.1720.1618.3055.00;
        }
    }
}

user@PE1# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    label-switched-path toPE2Primary192.0.2.6 {
        to 192.0.2.6;
        egress-protection;
    }
}

```

```

    interface all;
  }
  bgp {
    group ibgp {
      type internal;
      local-address 172.16.183.55;
      family inet-vpn {
        unicast;
      }
      neighbor 172.16.183.56;
      neighbor 172.16.183.59;
    }
  }
  isis {
    level 1 disable;
    level 2 wide-metrics-only;
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0;
  }
  ldp {
    track-igp-metric;
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}

user@PE1# show routing-instances
vpn1 {
  instance-type vrf;
  interface ge-1/2/0.0;
  route-distinguisher 172.16.183.55:10;
  vrf-target target:10:10;
  routing-options {
    static {
      route 100.0.0.0/24 next-hop 10.1.0.1;
    }
  }
  protocols {
    ospf {
      area 0.0.0.0 {
        interface ge-1/2/0.0;
      }
    }
  }
}

user@PE1# show routing-options
autonomous-system 64510;

```

**Device PE2**

```

user@PE2# show interfaces
ge-1/2/0 {

```

```
unit 0 {
  description to_P1;
  family inet {
    address 10.4.0.2/24;
  }
  family iso;
  family mpls;
}
}
ge-1/2/1 {
  unit 0 {
    description to_PE3;
    family inet {
      address 10.7.0.1/24;
    }
    family iso;
    family mpls;
  }
}
ge-1/2/2 {
  unit 0 {
    description to_CE2;
    family inet {
      address 10.8.0.1/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 172.16.183.56/32;
    }
    family iso {
      address 49.0002.1720.1618.3056.00;
    }
  }
}
}

user@PE2# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  label-switched-path toPE1 {
    to 172.16.183.55;
  }
  label-switched-path toPE2Primary192.0.2.6 {
    to 192.0.2.6;
    egress-protection;
  }
  interface all;
  egress-protection {
    context-identifier 192.0.2.6 {
```

```

        primary;
        advertise-mode stub-proxy;
    }
}
}
bgp {
    group ibgp {
        type internal;
        local-address 172.16.183.56;
        family inet-vpn {
            unicast {
                egress-protection {
                    context-identifier {
                        192.0.2.6;
                    }
                }
            }
        }
    }
    neighbor 172.16.183.55;
    neighbor 172.16.183.59;
}
}
isis {
    level 1 disable;
    level 2 wide-metrics-only;
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
ldp {
    track-igp-metric;
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}

user@PE2# show routing-options
autonomous-system 64510;

```

### Device PE3

```

user@PE3# show interfaces
ge-1/2/0 {
    unit 0 {
        description to_CE2;
        family inet {
            address 10.9.0.1/24;
        }
    }
}
ge-1/2/1 {
    unit 0 {
        description to_PE2;
    }
}

```

```
        family inet {
            address 10.7.0.2/24;
        }
        family iso;
        family mpls;
    }
}
ge-1/2/2 {
    unit 0 {
        description to_P1;
        family inet {
            address 10.5.0.2/24;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.183.59/32;
        }
        family iso {
            address 49.0002.1720.1618.3059.00;
        }
    }
}

user@PE3# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    label-switched-path toPE1 {
        to 172.16.183.55;
    }
    interface all;
    egress-protection {
        context-identifier 192.0.2.6 {
            protector;
            advertise-mode stub-proxy;
        }
    }
}
bgp {
    group ibgp {
        type internal;
        local-address 172.16.183.59;
        family inet-vpn {
            unicast {
                egress-protection {
                    keep-import remote-vrf;
                }
            }
        }
    }
}
```

```

    }
  }
  neighbor 172.16.183.55;
  neighbor 172.16.183.56;
}
}
isis {
  level 1 disable;
  level 2 wide-metrics-only;
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0;
}
ldp {
  track-igp-metric;
  interface all;
}

user@PE3# show policy-options
policy-statement remote-vrf {
  from community [ rsite1 rsite24 ];
  then accept;
}
community rsite1 members target:1:1;
community rsite24 members target:100:1023;

user@PE3# show routing-options
autonomous-system 64510;

```

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

## Verification

Confirm that the configuration is working properly.

- [Verifying the Protector Node on page 859](#)
- [Verifying the Primary Node on page 860](#)
- [Checking the Context Identifier Route on page 860](#)
- [Verifying Egress Protection on page 861](#)
- [Verifying the Routing Instance on Device PE1 on page 862](#)
- [Verifying the LSPs on page 862](#)
- [Verifying BGP NRLI on page 867](#)
- [Verifying the Traffic Engineering Database on page 869](#)
- [Verifying the IS-IS Database on page 874](#)

### Verifying the Protector Node

**Purpose** On the protector node (Device PE3), check the information about configured egress protection context identifiers.

**Action** user@PE3> show mpls context-identifier detail protector

```
ID: 192.0.2.6
Type: protector, Metric: 16777215, Mode: proxy
Context table: __PE3:192.0.2.6__.mpls.0
Context LSPs:
  toPE2Primary192.0.2.6, from: 172.16.183.55
  toPE2Primary192.0.2.6, from: 172.16.183.56

Total 1, Primary 0, Protector 1
```

**Meaning** Device PE3 is the protector node for two LSPs configured from Device PE1 (172.16.183.55) and Device PE2 (172.16.183.56).

---

### Verifying the Primary Node

**Purpose** On the primary node (Device PE2), check the information about configured egress protection context identifiers.

**Action** user@PE2> show mpls context-identifier detail primary

```
ID: 192.0.2.6
Type: primary, Metric: 1, Mode: proxy

Total 1, Primary 1, Protector 0
```

**Meaning** Device PE2 is the primary node.

---

### Checking the Context Identifier Route

**Purpose** Examine the information about the context identifier (192.0.2.6).



**Action** user@PE1> show route 192.0.2.6

```
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.6/24          *[IS-IS/18] 00:53:39, metric 21
> to 10.2.0.2 via ge-1/2/1.0
```

```
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.6/24          *[LDP/9] 00:53:39, metric 21
> to 10.2.0.2 via ge-1/2/1.0, Push 299808
```

user@PE2> show route 192.0.2.6

```
inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.6/24          *[MPLS/1] 3d 02:53:37, metric 1
Receive
[IS-IS/18] 00:06:08, metric 16777224
> to 10.7.0.2 via ge-1/2/1.0
```

user@PE3> show route 192.0.2.6

```
inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.6/24          *[MPLS/2] 3d 02:53:36, metric 16777215
Receive
[IS-IS/18] 3d 02:53:28, metric 11
> to 10.7.0.1 via ge-1/2/1.0
```

user@P1> show route 192.0.2.6

```
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.6/24          *[IS-IS/18] 00:53:40, metric 11
> to 10.4.0.2 via ge-1/2/0.0
```

```
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.0.2.6/24          *[LDP/9] 00:53:40, metric 11
> to 10.4.0.2 via ge-1/2/0.0
```

### Verifying Egress Protection

**Purpose** On Device PE3, check the routes in the routing table.

**Action** user@PE3> show mpls egress-protection detail

Instance	Type	Protection-Type
rsite1	remote-vrf	Protector
Route Target 1:1		
rsite24	remote-vrf	Protector
Route Target 100:1023		

**Meaning** **Instance** indicates the community name. **Type** shows the type of the VRF. It can be either **local-vrf** or **remote-vrf**. **Route Target** shows the route target associated with the routing instance.

---

### Verifying the Routing Instance on Device PE1

**Purpose** On Device PE1, check the routes in the routing table.

**Action** user@PE1> show route instance vpn1 detail

vpn1:

Router ID: 10.1.0.2

Type: vrf State: Active

Interfaces:

ge-1/2/0.0

Route-distinguisher: 172.16.183.55:10

Vrf-import: [ \_\_vrf-import-vpn1-internal\_\_ ]

Vrf-export: [ \_\_vrf-export-vpn1-internal\_\_ ]

Vrf-import-target: [ target:10:10 ]

Vrf-export-target: [ target:10:10 ]

Fast-reroute-priority: low

Tables:

vpn1.inet.0 : 4 routes (4 active, 0 holddown, 0 hidden)

---

### Verifying the LSPs

**Purpose** On all devices, check the LSP information.

**Action** user@PE1> show mpls lsp extensive

Ingress LSP: 1 sessions

192.0.2.6

From: 172.16.183.55, State: Up, ActiveRoute: 0, LSPname: toPE2Primary192.0.2.6

ActivePath: (primary)

LSPtype: Static Configured, Penultimate hop popping

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

\*Primary State: Up

Priorities: 7 0

SmartOptimizeTimer: 180

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 16777234)

10.2.0.2 S 10.5.0.2 S 192.0.2.6 S (link-id=2)

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):

10.2.0.2 10.5.0.2

17 Jun 10 13:13:04.973 CSPF: computation result accepted 10.2.0.2 10.5.0.2

192.0.2.6(link-id=2)

16 Jun 10 13:12:36.155 CSPF failed: no route toward 192.0.2.6[4 times]

15 Jun 10 13:11:26.269 CSPF: link down/deleted:

0.0.0.0(172.16.183.59:2147618818)(PE3.00/172.16.183.59)->0.0.0.0(192.0.2.6:2)(PE2-192.0.2.6.00/192.0.2.6)

14 Jun 10 13:10:11.771 Selected as active path

13 Jun 10 13:10:11.770 Record Route: 10.2.0.2 10.5.0.2

12 Jun 10 13:10:11.770 Up

11 Jun 10 13:10:11.634 Originate Call

10 Jun 10 13:10:11.634 CSPF: computation result accepted 10.2.0.2 10.5.0.2

192.0.2.6(link-id=2)

9 Jun 10 13:10:11.623 Clear Call

8 Jun 10 13:10:11.622 Deselected as active

7 Jun 7 11:23:08.224 Selected as active path

6 Jun 7 11:23:08.224 Record Route: 10.2.0.2 10.5.0.2

5 Jun 7 11:23:08.223 Up

4 Jun 7 11:23:08.116 Originate Call

3 Jun 7 11:23:08.116 CSPF: computation result accepted 10.2.0.2 10.5.0.2

192.0.2.6(link-id=2)

2 Jun 7 11:22:38.132 CSPF failed: no route toward 192.0.2.6

1 Jun 7 11:22:08.607 CSPF: could not determine self[8 times]

Created: Fri Jun 7 11:18:46 2013

Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions

172.16.183.55

From: 172.16.183.59, LSPstate: Up, ActiveRoute: 0

LSPname: toPE1, LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: -

Resv style: 1 FF, Label in: 3, Label out: -

Time left: 126, Since: Mon Jun 10 13:10:11 2013

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 2 receiver 10941 protocol 0

PATH rcvfrom: 10.2.0.2 (ge-1/2/1.0) 105 pkts

Adspec: received MTU 1500

PATH sentto: localclient

RESV rcvfrom: localclient

Record route: 10.5.0.2 10.2.0.2 <self>

```

172.16.183.55
  From: 172.16.183.56, LSPstate: Up, ActiveRoute: 0
  LSPname: toPE1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 156, Since: Mon Jun 10 13:10:11 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 59956 protocol 0
  PATH rcvfrom: 10.2.0.2 (ge-1/2/1.0) 105 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.4.0.2 10.2.0.2 <self>
Total 2 displayed, Up 2, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
-----

user@PE2> show mpls lsp extensive
Ingress LSP: 2 sessions

192.0.2.6
  From: 172.16.183.56, State: Up, ActiveRoute: 0, LSPname: toPE2Primary192.0.2.6

  ActivePath: (primary)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 16777224)
10.7.0.2 S 192.0.2.6 S (link-id=2)
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
      10.7.0.2
16 Jun 10 13:13:07.220 CSPF: computation result accepted 10.7.0.2
192.0.2.6(link-id=2)
15 Jun 10 13:12:38.250 CSPF failed: no route toward 192.0.2.6[4 times]
14 Jun 10 13:11:26.258 CSPF: link down/deleted:
0.0.0.0(172.16.183.59:2147618818)(PE3.00/172.16.183.59)->0.0.0.0(192.0.2.6:2)(PE2-192.0.2.6.00/192.0.2.6)

13 Jun 10 13:10:11.746 Selected as active path
12 Jun 10 13:10:11.743 Record Route: 10.7.0.2
11 Jun 10 13:10:11.742 Up
10 Jun 10 13:10:11.680 Originate Call
9 Jun 10 13:10:11.680 CSPF: computation result accepted 10.7.0.2
192.0.2.6(link-id=2)
8 Jun 10 13:10:11.674 Clear Call
7 Jun 10 13:10:11.669 Deselected as active
6 Jun 7 11:23:09.370 Selected as active path
5 Jun 7 11:23:09.370 Record Route: 10.7.0.2
4 Jun 7 11:23:09.369 Up
3 Jun 7 11:23:09.349 Originate Call
2 Jun 7 11:23:09.349 CSPF: computation result accepted 10.7.0.2
192.0.2.6(link-id=2)
1 Jun 7 11:22:40.140 CSPF failed: no route toward 192.0.2.6[9 times]
Created: Fri Jun 7 11:18:46 2013

```

```

172.16.183.55
  From: 172.16.183.56, State: Up, ActiveRoute: 0, LSPName: toPE1
  ActivePath: (primary)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.4.0.1 S 10.2.0.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.4.0.1 10.2.0.1
13 Jun 10 13:10:11.794 Selected as active path
12 Jun 10 13:10:11.793 Record Route: 10.4.0.1 10.2.0.1
11 Jun 10 13:10:11.793 Up
10 Jun 10 13:10:11.679 Originate Call
9 Jun 10 13:10:11.679 CSPF: computation result accepted 10.4.0.1 10.2.0.1
8 Jun 10 13:10:11.660 Clear Call
7 Jun 10 13:10:11.645 Deselected as active
6 Jun 7 11:22:40.031 Selected as active path
5 Jun 7 11:22:40.024 Record Route: 10.4.0.1 10.2.0.1
4 Jun 7 11:22:40.012 Up
3 Jun 7 11:22:39.687 Originate Call
2 Jun 7 11:22:39.687 CSPF: computation result accepted 10.4.0.1 10.2.0.1
1 Jun 7 11:22:10.235 CSPF failed: no route toward 172.16.183.55[8 times]
  Created: Fri Jun 7 11:18:45 2013
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@PE3> show mpls lsp extensive
Ingress LSP: 1 sessions

```

```

172.16.183.55
  From: 172.16.183.59, State: Up, ActiveRoute: 0, LSPName: toPE1
  ActivePath: (primary)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.5.0.1 S 10.2.0.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.5.0.1 10.2.0.1
13 Jun 10 13:10:11.708 Selected as active path
12 Jun 10 13:10:11.703 Record Route: 10.5.0.1 10.2.0.1
11 Jun 10 13:10:11.703 Up
10 Jun 10 13:10:11.599 Originate Call
9 Jun 10 13:10:11.599 CSPF: computation result accepted 10.5.0.1 10.2.0.1
8 Jun 10 13:10:11.558 Clear Call
7 Jun 10 13:10:11.555 Deselected as active
6 Jun 7 11:22:41.829 Selected as active path
5 Jun 7 11:22:41.828 Record Route: 10.5.0.1 10.2.0.1

```

```

4 Jun  7 11:22:41.827 Up
3 Jun  7 11:22:41.767 Originate Call
2 Jun  7 11:22:41.767 CSPF: computation result accepted 10.5.0.1 10.2.0.1
1 Jun  7 11:22:12.289 CSPF failed: no route toward 172.16.183.55[8 times]
Created: Fri Jun  7 11:18:45 2013
Total 1 displayed, Up 1, Down 0

```

Egress LSP: 2 sessions

192.0.2.6

```

From: 172.16.183.55, LSPstate: Up, ActiveRoute: 0
LSPname: toPE2Primary192.0.2.6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 299920, Label out: 3
Time left: 141, Since: Mon Jun 10 13:10:11 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 17060 protocol 0
Attrib flags: Non-PHP OOB
PATH rcvfrom: 10.5.0.1 (ge-1/2/2.0) 105 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.2.0.1 10.5.0.1 <self>

```

192.0.2.6

```

From: 172.16.183.56, LSPstate: Up, ActiveRoute: 0
LSPname: toPE2Primary192.0.2.6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 299936, Label out: 3
Time left: 152, Since: Mon Jun 10 13:10:11 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 59957 protocol 0
Attrib flags: Non-PHP OOB
PATH rcvfrom: 10.7.0.1 (ge-1/2/1.0) 106 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.7.0.1 <self>

```

Total 2 displayed, Up 2, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

user@P1> show mpls lsp extensive

Ingress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit LSP: 3 sessions

192.0.2.6

```

From: 172.16.183.55, LSPstate: Up, ActiveRoute: 0
LSPname: toPE2Primary192.0.2.6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299920
Resv style: 1 FF, Label in: 299904, Label out: 299920
Time left: 141, Since: Mon Jun 10 13:10:11 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

```

```

Port number: sender 2 receiver 17060 protocol 0
Attrib flags: Non-PHP OOB
PATH rcvfrom: 10.2.0.1 (ge-1/2/1.0) 106 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.5.0.2 (ge-1/2/2.0) 105 pkts
RESV rcvfrom: 10.5.0.2 (ge-1/2/2.0) 105 pkts
Explct route: 10.5.0.2 192.0.2.6 (link-id=2)
Record route: 10.2.0.1 <self> 10.5.0.2

172.16.183.55
From: 172.16.183.59, LSPstate: Up, ActiveRoute: 0
LSPname: toPE1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 299888, Label out: 3
Time left: 158, Since: Mon Jun 10 13:10:11 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 10941 protocol 0
PATH rcvfrom: 10.5.0.2 (ge-1/2/2.0) 106 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.2.0.1 (ge-1/2/1.0) 105 pkts
RESV rcvfrom: 10.2.0.1 (ge-1/2/1.0) 105 pkts
Explct route: 10.2.0.1
Record route: 10.5.0.2 <self> 10.2.0.1

172.16.183.55
From: 172.16.183.56, LSPstate: Up, ActiveRoute: 0
LSPname: toPE1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 299920, Label out: 3
Time left: 141, Since: Mon Jun 10 13:10:11 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 59956 protocol 0
PATH rcvfrom: 10.4.0.2 (ge-1/2/0.0) 105 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.2.0.1 (ge-1/2/1.0) 105 pkts
RESV rcvfrom: 10.2.0.1 (ge-1/2/1.0) 105 pkts
Explct route: 10.2.0.1
Record route: 10.4.0.2 <self> 10.2.0.1
Total 3 displayed, Up 3, Down 0

```

### Verifying BGP NRLI

**Purpose** Check the details of the BGP VPN network layer reachability information.

**Action** user@PE3> show bgp neighbor

```

Peer: 172.16.183.55+179 AS 64510 Local: 172.16.183.59+61747 AS 64510
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 172.16.183.59 Holdtime: 90 Preference: 170
NLRI configured with egress-protection: inet-vpn-unicast
Egress-protection NLRI inet-vpn-unicast, keep-import: [ remote-vrf ]
  Number of flaps: 0
  Peer ID: 172.16.183.55    Local ID: 172.16.183.59    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  Peer supports 4 byte AS extension (peer-as 64510)
  Peer does not support Addpath
  Table bgp.13vpn.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:          0
    Received prefixes:        0
    Accepted prefixes:        0
    Suppressed due to damping: 0
  Last traffic (seconds): Received 25    Sent 21    Checked 11
  Input messages: Total 32046 Updates 7    Refreshes 0    Octets 609365
  Output messages: Total 32050 Updates 0    Refreshes 5    Octets 609010
  Output Queue[0]: 0

Peer: 172.16.183.56+62754 AS 64510 Local: 172.16.183.59+179 AS 64510
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 172.16.183.59 Holdtime: 90 Preference: 170
NLRI configured with egress-protection: inet-vpn-unicast
Egress-protection NLRI inet-vpn-unicast, keep-import: [ remote-vrf ]
  Number of flaps: 1
  Last flap event: TransportError
  Peer ID: 172.16.183.56    Local ID: 172.16.183.59    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 1
  BFD: disabled, down
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-vpn-unicast
  Peer supports 4 byte AS extension (peer-as 64510)

```



```
Peer does not support Addpath
Table bgp.l3vpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 19   Sent 8   Checked 34
  Input messages:  Total 10025  Updates 0   Refreshes 2   Octets 190523
  Output messages: Total 10024  Updates 0   Refreshes 2   Octets 190504
  Output Queue[0]: 0
```

**Meaning** NLRI configured with `egress-protection` shows the BGP family configured with egress protection. `egress-protection NLRI inet-vpn-unicast, keep-import: [remote-vrf]` shows the egress protection routing policy for the BGP group.

### Verifying the Traffic Engineering Database

---

**Purpose** On all devices, check the TED.

**Action** user@PE1> show ted database

```

TED database: 9 ISIS nodes 5 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.00(172.16.0.3)                Rtr   44      3      3 IS-IS(2)
  To: P1.02, Local: 10.2.0.2, Remote: 0.0.0.0
    Local interface index: 149, Remote interface index: 0
  To: PE2.02, Local: 10.4.0.1, Remote: 0.0.0.0
    Local interface index: 150, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.1, Remote: 0.0.0.0
    Local interface index: 133, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.02                             Net   111      2      2 IS-IS(2)
  To: PE1.00(172.16.183.55), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2-192.0.2.6.00(192.0.2.6)      Rtr   345      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 1, Remote interface index: 2147618817
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2, Remote interface index: 2147618818
ID                               Type Age(s) LnkIn LnkOut Protocol
PE1.00(172.16.183.55)            Rtr   487      1      1 IS-IS(2)
  To: P1.02, Local: 10.2.0.1, Remote: 0.0.0.0
    Local interface index: 148, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.00(172.16.183.56)            Rtr   353      3      3 IS-IS(2)
  To: PE2.02, Local: 10.4.0.2, Remote: 0.0.0.0
    Local interface index: 155, Remote interface index: 0
  To: PE3.02, Local: 10.7.0.1, Remote: 0.0.0.0
    Local interface index: 153, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2147618817, Remote interface index: 1
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.02                             Net    59      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.00(172.16.183.59)            Rtr   435      3      3 IS-IS(2)
  To: PE3.02, Local: 10.7.0.2, Remote: 0.0.0.0
    Local interface index: 154, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.2, Remote: 0.0.0.0
    Local interface index: 158, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2147618818, Remote interface index: 2
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.02                             Net   706      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.03                             Net   583      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0

```

```

Local interface index: 0, Remote interface index: 0

user@PE2> show ted database
TED database: 9 ISIS nodes 5 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.00(172.16.0.3)                Rtr   44      3      3 IS-IS(2)
  To: PE2.02, Local: 10.4.0.1, Remote: 0.0.0.0
  Local interface index: 150, Remote interface index: 0
  To: P1.02, Local: 10.2.0.2, Remote: 0.0.0.0
  Local interface index: 149, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.1, Remote: 0.0.0.0
  Local interface index: 133, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.02                             Net  111      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: PE1.00(172.16.183.55), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2-192.0.2.6.00(192.0.2.6)      Rtr  345      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 1, Remote interface index: 2147618817
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 2, Remote interface index: 2147618818
ID                               Type Age(s) LnkIn LnkOut Protocol
PE1.00(172.16.183.55)            Rtr  487      1      1 IS-IS(2)
  To: P1.02, Local: 10.2.0.1, Remote: 0.0.0.0
  Local interface index: 148, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.00(172.16.183.56)            Rtr  353      3      3 IS-IS(2)
  To: PE2.02, Local: 10.4.0.2, Remote: 0.0.0.0
  Local interface index: 155, Remote interface index: 0
  To: PE3.02, Local: 10.7.0.1, Remote: 0.0.0.0
  Local interface index: 153, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 2147618817, Remote interface index: 1
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.02                             Net   60      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.00(172.16.183.59)            Rtr  435      3      3 IS-IS(2)
  To: PE3.02, Local: 10.7.0.2, Remote: 0.0.0.0
  Local interface index: 154, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.2, Remote: 0.0.0.0
  Local interface index: 158, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 2147618818, Remote interface index: 2
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.02                             Net  706      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.03                             Net  583      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0

```

```

user@PE3> show ted database
TED database: 9 ISIS nodes 5 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.00(172.16.0.3)                Rtr    44      3      3 IS-IS(2)
  To: P1.02, Local: 10.2.0.2, Remote: 0.0.0.0
    Local interface index: 149, Remote interface index: 0
  To: PE2.02, Local: 10.4.0.1, Remote: 0.0.0.0
    Local interface index: 150, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.1, Remote: 0.0.0.0
    Local interface index: 133, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.02                            Net    111      2      2 IS-IS(2)
  To: PE1.00(172.16.183.55), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2-192.0.2.6.00(192.0.2.6)      Rtr   345      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 1, Remote interface index: 2147618817
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2, Remote interface index: 2147618818
ID                               Type Age(s) LnkIn LnkOut Protocol
PE1.00(172.16.183.55)            Rtr   487      1      1 IS-IS(2)
  To: P1.02, Local: 10.2.0.1, Remote: 0.0.0.0
    Local interface index: 148, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.00(172.16.183.56)            Rtr   353      3      3 IS-IS(2)
  To: PE3.02, Local: 10.7.0.1, Remote: 0.0.0.0
    Local interface index: 153, Remote interface index: 0
  To: PE2.02, Local: 10.4.0.2, Remote: 0.0.0.0
    Local interface index: 155, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2147618817, Remote interface index: 1
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.02                            Net    59      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.00(172.16.183.59)            Rtr   435      3      3 IS-IS(2)
  To: PE3.02, Local: 10.7.0.2, Remote: 0.0.0.0
    Local interface index: 154, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.2, Remote: 0.0.0.0
    Local interface index: 158, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2147618818, Remote interface index: 2
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.02                            Net   706      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.03                            Net   583      2      2 IS-IS(2)
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
-----

```

```
user@P1> show ted database
```

```
TED database: 9 ISIS nodes 5 INET nodes
```

```
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.00(172.16.0.3)                Rtr   44      3      3 IS-IS(2)
  To: PE2.02, Local: 10.4.0.1, Remote: 0.0.0.0
    Local interface index: 150, Remote interface index: 0
  To: P1.02, Local: 10.2.0.2, Remote: 0.0.0.0
    Local interface index: 149, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.1, Remote: 0.0.0.0
    Local interface index: 133, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
P1.02                             Net   111      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE1.00(172.16.183.55), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2-192.0.2.6.00(192.0.2.6)      Rtr  345      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 1, Remote interface index: 2147618817
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2, Remote interface index: 2147618818
ID                               Type Age(s) LnkIn LnkOut Protocol
PE1.00(172.16.183.55)            Rtr   487      1      1 IS-IS(2)
  To: P1.02, Local: 10.2.0.1, Remote: 0.0.0.0
    Local interface index: 148, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.00(172.16.183.56)            Rtr   353      3      3 IS-IS(2)
  To: PE2.02, Local: 10.4.0.2, Remote: 0.0.0.0
    Local interface index: 155, Remote interface index: 0
  To: PE3.02, Local: 10.7.0.1, Remote: 0.0.0.0
    Local interface index: 153, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2147618817, Remote interface index: 1
ID                               Type Age(s) LnkIn LnkOut Protocol
PE2.02                             Net    59      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.00(172.16.183.59)            Rtr   435      3      3 IS-IS(2)
  To: PE3.02, Local: 10.7.0.2, Remote: 0.0.0.0
    Local interface index: 154, Remote interface index: 0
  To: PE3.03, Local: 10.5.0.2, Remote: 0.0.0.0
    Local interface index: 158, Remote interface index: 0
  To: PE2-192.0.2.6.00(192.0.2.6), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 2147618818, Remote interface index: 2
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.02                             Net   706      2      2 IS-IS(2)
  To: PE2.00(172.16.183.56), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
PE3.03                             Net   583      2      2 IS-IS(2)
  To: P1.00(172.16.0.3), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: PE3.00(172.16.183.59), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
```

### Verifying the IS-IS Database

---

**Purpose** On all devices, check the IS-IS database.

**Action** user@PE1> show isis database

IS-IS level 1 link-state database:  
0 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
P1.00-00	0x46b	0x1924	590	L1 L2
P1.02-00	0x465	0xe67a	523	L1 L2
PE2-192.0.2.6.00-00	0xd0e	0x6b8d	1086	L1 L2 Overload
PE1.00-00	0x46f	0xa8b	992	L1 L2
PE2.00-00	0x46b	0xefd6	1077	L1 L2
PE2.02-00	0x464	0x4db4	573	L1 L2
PE3.00-00	0x46f	0xb6e8	1016	L1 L2
PE3.02-00	0x465	0x2675	762	L1 L2
PE3.03-00	0x465	0x47b2	797	L1 L2

9 LSPs

user@PE2> show isis database

IS-IS level 1 link-state database:  
0 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
P1.00-00	0x46b	0x1924	590	L1 L2
P1.02-00	0x465	0xe67a	523	L1 L2
PE2-192.0.2.6.00-00	0xd0e	0x6b8d	1090	L1 L2 Overload
PE1.00-00	0x46f	0xa8b	988	L1 L2
PE2.00-00	0x46b	0xefd6	1080	L1 L2
PE2.02-00	0x464	0x4db4	576	L1 L2
PE3.00-00	0x46f	0xb6e8	1018	L1 L2
PE3.02-00	0x465	0x2675	763	L1 L2
PE3.03-00	0x465	0x47b2	799	L1 L2

9 LSPs

user@PE3> show isis database

IS-IS level 1 link-state database:  
0 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
P1.00-00	0x46b	0x1924	590	L1 L2
P1.02-00	0x465	0xe67a	523	L1 L2
PE2-192.0.2.6.00-00	0xd0e	0x6b8d	1088	L1 L2 Overload
PE1.00-00	0x46f	0xa8b	988	L1 L2
PE2.00-00	0x46b	0xefd6	1079	L1 L2
PE2.02-00	0x464	0x4db4	575	L1 L2
PE3.00-00	0x46f	0xb6e8	1020	L1 L2
PE3.02-00	0x465	0x2675	765	L1 L2
PE3.03-00	0x465	0x47b2	801	L1 L2

9 LSPs

user@P1> show isis database

IS-IS level 1 link-state database:  
0 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
P1.00-00	0x46b	0x1924	592	L1 L2
P1.02-00	0x465	0xe67a	525	L1 L2
PE2-192.0.2.6.00-00	0xd0e	0x6b8d	1088	L1 L2 Overload
PE1.00-00	0x46f	0xa8b	990	L1 L2

PE2.00-00	0x46b	0xefd6	1079	L1	L2
PE2.02-00	0x464	0x4db4	575	L1	L2
PE3.00-00	0x46f	0xb6e8	1018	L1	L2
PE3.02-00	0x465	0x2675	763	L1	L2
PE3.03-00	0x465	0x47b2	799	L1	L2
9 LSPs					

- Related Documentation**
- <http://www.ethernetacademy.net/Ethernet-Academy-Articles/end-to-end-service-restoration>
  - [Egress Protection for Layer 3 VPN Edge Protection Overview on page 827](#)



# Configuring Provider Edge Link Protection in Layer 3 VPNs

- [Understanding Provider Edge Link Protection for BGP Labeled Unicast Paths on page 877](#)
- [Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 878](#)
- [Example: Configuring Provider Edge Link Protection for BGP Labeled Unicast Paths on page 895](#)

## Understanding Provider Edge Link Protection for BGP Labeled Unicast Paths

---

In MPLS service provider networks, when Layer 3 VPNs are used for carrier-of-carriers deployments, the protocol used to link the customer edge (CE) routers in one autonomous system (AS) and a provider edge (PE) router in another AS is BGP labeled-unicast. Reroute solutions between ASs are essential to help service providers ensure that network outages will have minimal impact on the data flows through the networks. A service provider that is a customer of another service provider can have different CE routers that are connected to the other service provider through different PE routers. This setup enables load balancing of traffic. However, this can lead to disruption in traffic if the link between one CE router and a PE router goes down. Therefore, a precomputed protection path should be configured such that if a link between a CE router and a PE router goes down, the protection path (also known as the backup path) between the other CE router and the alternative PE router can be used.

To configure a labeled-unicast path to be a protection path, use the **protection** statement at the `[edit routing-instances instance-name protocols bgp family inet labeled-unicast]` hierarchy level:

```
routing-instances {
  customer {
    instance-type vrf;
    ...
    protocols {
      bgp {
        family inet {
          labeled-unicast {
            protection;
          }
        }
        family inet6 {
```

```

        labeled-unicast {
            protection;
        }
    }
    type external;
    ...
}
}
}
}

```

The **protection** statement indicates that protection is desired on prefixes received from the particular neighbor or family. After protection is enabled for a given family, group, or neighbor, protection entries are added for prefixes or next hops received from the given peer.



**NOTE:** A protection path can be selected only if the best path has already been installed by BGP in the forwarding table. This is because a protection path cannot be used as the best path.

To minimize packet loss when the protected path is down, also use the **per-prefix-label** statement at the [edit routing-instances *instance-name* protocols **bgp** family inet labeled-unicast] hierarchy level. Set this statement on every PE router within the AS containing the protected path.

The protection path selection takes place based on the value of two state flags:

- The **ProtectionPath** flag indicates paths desiring protection.
- The **ProtectionCand** flag indicates the route entry that can be used as a protection path.



**NOTE:**

- Provider edge link protection is configured only for external peers.
- If provider edge link protection is configured with the **equal-external-internal** multipath statement, multipath takes precedence over protection.

#### Related Documentation

- [Example: Configuring Provider Edge Link Protection for BGP Labeled Unicast Paths on page 895](#)

## Example: Configuring Provider Edge Link Protection in Layer 3 VPNs

- [Understanding Provider Edge Link Protection in Layer 3 VPNs on page 879](#)
- [Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 880](#)

## Understanding Provider Edge Link Protection in Layer 3 VPNs

In an MPLS service provider network, a customer can have dual-homed CE routers that are connected to the service provider through different PE routers. This setup enables load balancing of traffic in the service provider network. However, this can lead to disruption in traffic if the link between a CE router and a PE router goes down. Hence, a precomputed protection path should be configured such that if a link between a CE router and a PE router goes down, the protection path (also known as the backup path) between the CE router and an alternate PE router can be used.

To configure a path to be a protection path, use the **protection** statement at the **[edit routing-instances *instance-name* protocols bgp family inet unicast]** hierarchy level:

```
routing-instances {
  customer {
    instance-type vrf;
    ...
    protocols {
      bgp {
        type external;
        ...
        family inet {
          unicast {
            protection;
          }
        }
        family inet6 {
          unicast {
            protection;
          }
        }
      }
    }
  }
}
```

The **protection** statement indicates that protection is required on prefixes received from the particular neighbor or family. After protection is enabled for a given family, group, or neighbor, protection entries are added for prefixes or next hops received from the given peer.



**NOTE:** A protection path can be selected only if the best path has already been installed by BGP in the forwarding table. This is because a protection path cannot be used as the best path.



**NOTE:** The option **vrf-table-label** must be configured under the **[routing-instances *instance-name*]** hierarchy for the routers that have protected PE-CE links. This applies to Junos OS Releases 12.3 through 13.2 inclusive.

The protection path selection takes place based on the value of two state flags:

- The **ProtectionPath** flag indicates paths requesting protection.
- The **ProtectionCand** flag indicates the route entry that can be used as a protection path.



NOTE:

- Provider edge link protection is configured only for external peers.
  - If provider edge link protection is configured with the **equal-external-internal** multipath statement, multipath takes precedence over protection.
- 

**See Also** • [Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 880](#)

## Example: Configuring Provider Edge Link Protection in Layer 3 VPNs

This example shows how to configure a provider edge protection path that can be used in case of a link failure in an MPLS network.

- [Requirements on page 880](#)
- [Overview on page 880](#)
- [Configuration on page 881](#)
- [Verification on page 891](#)

### Requirements

---

This example uses the following hardware components, software components and configuration options:

- M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers
- Junos OS Release 12.3 through 13.2 inclusive
- The option **vrf-table-label** must be enabled at the **[routing-instances *instance-name*]** hierarchy level for routers with protected PE-CE links.

### Overview

---

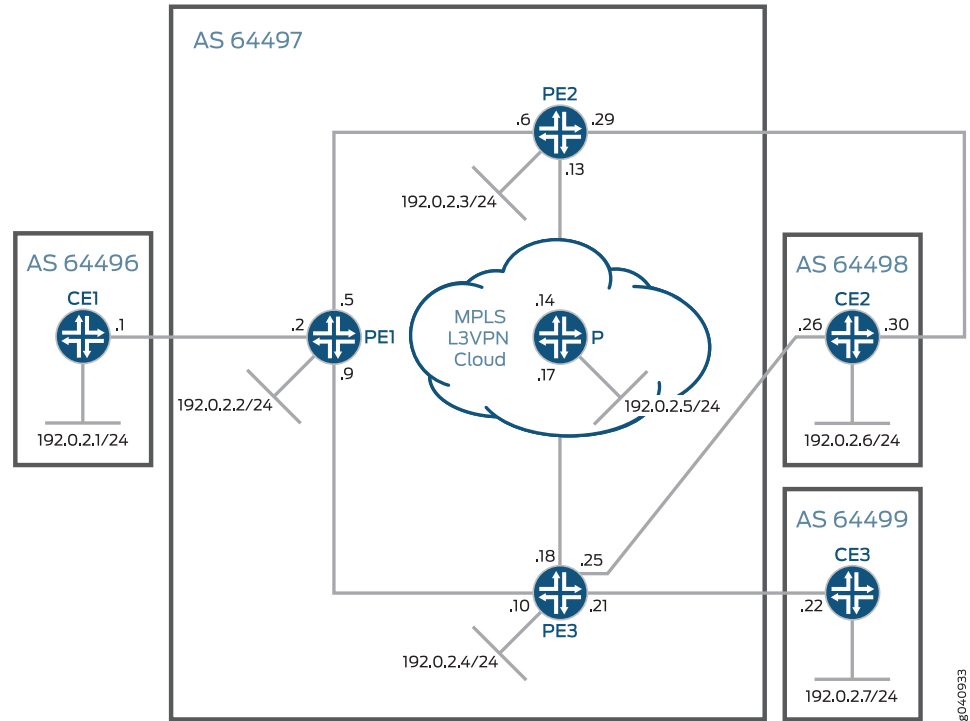
The following example shows how to configure provider edge link protection in a Layer 3 VPN.

### Topology

In this example, a Layer 3 VPN is set up by configuring three customer edge devices and three service provider edge devices in four autonomous systems. The CE devices are configured in AS 64496, AS 64498, and AS 64499. The PE devices are configured in AS 64497.

Figure 82 on page 881 shows the topology used in this example.

Figure 82: Provider Edge Link Protection in a Layer 3 VPN



The aim of this example is to protect the provider edge link between Routers PE2 and CE2. Protection is configured on the backup link between Routers PE3 and CE2, such that the traffic can be routed through this link when the PE2-CE2 link goes down.

### 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.

#### Router CE1

```
set interfaces ge-2/0/0 unit 0 description toPE1
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.1/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.1/24
set interfaces lo0 unit 0 family inet6 address 2001:db8::1/128
set routing-options router-id 192.0.2.1
set routing-options autonomous-system 64496
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 export send-direct
set protocols bgp group toPE1 peer-as 64497
set protocols bgp group toPE1 neighbor 10.1.1.2
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

```

Router PE1
set interfaces ge-2/0/0 unit 0 description toCE1
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.2/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toPE2
set interfaces ge-2/0/1 unit 0 family inet address 10.1.1.5/30
set interfaces ge-2/0/1 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 description toPE3
set interfaces ge-2/0/2 unit 0 family inet address 10.1.1.9/30
set interfaces ge-2/0/2 unit 0 family inet6 address 2001:db8:0:9::/64 eui-64
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.2/24
set interfaces lo0 unit 0 family inet6 address 2001:db8::2/128
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0 metric 10
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-2/0/1.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-2/0/2.0 metric 10
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal family inet6-vpn unicast
set protocols bgp group toInternal multipath
set protocols bgp group toInternal local-address 192.0.2.2
set protocols bgp group toInternal neighbor 192.0.2.3
set protocols bgp group toInternal neighbor 192.0.2.4
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 64497
set routing-options forwarding-table export lb
set routing-instances radium instance-type vrf
set routing-instances radium interface ge-2/0/0.0
set routing-instances radium route-distinguisher 64497:1
set routing-instances radium vrf-target target:64497:1
set routing-instances radium protocols bgp group toCE1 type external
set routing-instances radium protocols bgp group toCE1 peer-as 64496
set routing-instances radium protocols bgp group toCE1 neighbor 10.1.1.1
set policy-options policy-statement lb then load-balance per-packet

```

```

Router PE2
set interfaces ge-2/0/0 unit 0 description toPE1
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.6/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toP
set interfaces ge-2/0/1 unit 0 family inet address 10.1.1.13/30
set interfaces ge-2/0/1 unit 0 family inet6 address 2001:db8:0:13::/64 eui-64
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 description toCE2
set interfaces ge-2/0/2 unit 0 family inet address 10.1.1.29/30
set interfaces ge-2/0/2 unit 0 family inet6 address 2001:db8:0:29::/64 eui-64
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.3/24
set interfaces lo0 unit 0 family inet6 address 2001:db8::3/128

```

```

set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 5
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface ge-2/0/1.0 metric 5
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal family inet6-vpn unicast
set protocols bgp group toInternal multipath
set protocols bgp group toInternal local-address 192.0.2.3
set protocols bgp group toInternal neighbor 192.0.2.2
set protocols bgp group toInternal neighbor 192.0.2.4
set routing-options router-id 192.0.2.3
set routing-options autonomous-system 64497
set routing-options forwarding-table export lb
set routing-instances radium instance-type vrf
set routing-instances radium interface ge-2/0/2.0
set routing-instances radium route-distinguisher 64497:1
set routing-instances radium vrf-target target:64497:1
set routing-instances radium protocols bgp group toCE2 type external
set routing-instances radium protocols bgp group toCE2 peer-as 64498
set routing-instances radium protocols bgp group toCE2 neighbor 10.1.1.30
set policy-options policy-statement lb then load-balance per-packet

```

#### Router PE3

```

set interfaces ge-2/0/0 unit 0 description toPE1
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.10/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:9::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toP
set interfaces ge-2/0/1 unit 0 family inet address 10.1.1.18/30
set interfaces ge-2/0/1 unit 0 family inet6 address 2001:db8:0:17::/64 eui-64
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 description toCE2
set interfaces ge-2/0/2 unit 0 family inet address 10.1.1.25/30
set interfaces ge-2/0/2 unit 0 family inet6 address 2001:db8:0:25::/64 eui-64
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces ge-2/0/3 unit 0 description toCE3
set interfaces ge-2/0/3 unit 0 family inet address 10.1.1.21/30
set interfaces ge-2/0/3 unit 0 family inet6 address 2001:db8:0:21::/64 eui-64
set interfaces ge-2/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.4/24
set interfaces lo0 unit 0 family inet6 address 2001:db8::4/128
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 5
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-2/0/1.0 metric 5
set protocols ospf3 area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast

```

```

set protocols bgp group toInternal family inet6-vpn unicast
set protocols bgp group toInternal multipath
set protocols bgp group toInternal local-address 192.0.2.4
set protocols bgp group toInternal neighbor 192.0.2.2
set protocols bgp group toInternal neighbor 192.0.2.3
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 64497
set routing-options forwarding-table export lb
set routing-instances radium instance-type vrf
set routing-instances radium vrf-table-label
set routing-instances radium interface ge-2/0/2.0
set routing-instances radium interface ge-2/0/3.0
set routing-instances radium route-distinguisher 64497:1
set routing-instances radium vrf-target target:64497:1
set routing-instances radium protocols bgp group toCE2 type external
set routing-instances radium protocols bgp group toCE2 peer-as 64498
set routing-instances radium protocols bgp group toCE2 neighbor 10.1.1.26
set routing-instances radium protocols bgp group toCE2 family inet unicast protection
set routing-instances radium protocols bgp group toCE2 family inet6 unicast protection
set routing-instances radium protocols bgp group toCE3 type external
set routing-instances radium protocols bgp group toCE3 peer-as 64499
set routing-instances radium protocols bgp group toCE3 neighbor 10.1.1.22
set policy-options policy-statement lb then load-balance per-packet

```

**Router P**

```

set interfaces ge-2/0/0 unit 0 description toPE2
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.14/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:13::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toPE3
set interfaces ge-2/0/1 unit 0 family inet address 10.1.1.17/30
set interfaces ge-2/0/1 unit 0 family inet6 address 2001:db8:0:17::/64 eui-64
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.5/24
set interfaces lo0 unit 0 family inet6 address 2001:db8::5/128
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 64497
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 5
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 5
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-2/0/0.0 metric 5
set protocols ospf3 area 0.0.0.0 interface ge-2/0/1.0 metric 5

```

**Router CE2**

```

set interfaces ge-2/0/0 unit 0 description toPE2
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.30/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:29::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toPE3
set interfaces ge-2/0/1 unit 0 family inet address 10.1.1.26/30
set interfaces ge-2/0/1 unit 0 family inet6 address 2001:db8:0:25::/64 eui-64
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.6/24

```



```

set interfaces lo0 unit 0 family inet6 address 2001:db8::6/128
set routing-options router-id 192.0.2.6
set routing-options autonomous-system 64498
set protocols bgp group toAS2 type external
set protocols bgp group toAS2 export send-direct
set protocols bgp group toAS2 peer-as 64497
set protocols bgp group toAS2 neighbor 10.1.1.25
set protocols bgp group toAS2 neighbor 10.1.1.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept

```

**Router CE3**

```

set interfaces ge-2/0/0 unit 0 description toPE3
set interfaces ge-2/0/0 unit 0 family inet address 10.1.1.22/30
set interfaces ge-2/0/0 unit 0 family inet6 address 2001:db8:0:21::/64 eui-64
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.7/24
set interfaces lo0 unit 0 family inet6 address 2001:db8::7/128
set routing-options router-id 192.0.2.7
set routing-options autonomous-system 64499
set protocols bgp group toPE3 type external
set protocols bgp group toPE3 export send-direct
set protocols bgp group toPE3 peer-as 64497
set protocols bgp group toPE3 neighbor 10.1.1.21
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept

```

### Configuring Provider Edge Link Protection in Layer 3 VPNs

**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 provider edge link protection:

1. Configure the router interfaces.

```

[edit interfaces]
user@PE3# set ge-2/0/0 unit 0 description toPE1
user@PE3# set ge-2/0/0 unit 0 family inet address 10.1.1.10/30
user@PE3# set ge-2/0/0 unit 0 family inet6 address 2001:db8:0:9::/64 eui-64
user@PE3# set ge-2/0/0 unit 0 family mpls

user@PE3# set ge-2/0/1 unit 0 description toP
user@PE3# set ge-2/0/1 unit 0 family inet address 10.1.1.18/30
user@PE3# set ge-2/0/1 unit 0 family inet6 address 2001:db8:0:17::/64 eui-64
user@PE3# set ge-2/0/1 unit 0 family mpls

user@PE3# set ge-2/0/2 unit 0 description toCE2
user@PE3# set ge-2/0/2 unit 0 family inet address 10.1.1.25/30
user@PE3# set ge-2/0/2 unit 0 family inet6 address 12001:db8:0:25::/64 eui-64
user@PE3# set ge-2/0/2 unit 0 family mpls

user@PE3# set ge-2/0/3 unit 0 description toCE3
user@PE3# set ge-2/0/3 unit 0 family inet address 10.1.1.21/30
user@PE3# set ge-2/0/3 unit 0 family inet6 address 2001:db8:0:21::/64 eui-64
user@PE3# set ge-2/0/3 unit 0 family mpls

```

```
user@PE3# set lo0 unit 0 family inet address 192.0.2.4/24
user@PE3# set lo0 unit 0 family inet6 address 2001:db8::4/128
```

Similarly, configure the interfaces on all other routers.

2. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@PE3# set router-id 192.0.2.4
user@PE3# set autonomous-system 64497
```

Similarly, configure the router ID and AS number for all other routers. In this example, the router ID is chosen to be identical to the loopback address configured on the router.

3. Configure MPLS and LDP on all interfaces of Router PE3.

```
[edit protocols]
user@PE3# set mpls interface all
user@PE3# set ldp interface all
```

Similarly, configure other PE routers.

4. Configure an IGP on the core-facing interfaces of Router PE3.

```
[edit protocols ospf area 0.0.0.0]
user@PE3# set interface lo0.0 passive
user@PE3# set interface ge-2/0/1.0 metric 5
user@PE3# set interface ge-2/0/0.0 metric 10

[edit protocols ospf3 area 0.0.0.0]
user@PE3# set interface lo0.0 passive
user@PE3# set interface ge-2/0/1.0 metric 5
user@PE3# set interface ge-2/0/0.0 metric 10
```

Similarly, configure other PE routers.

5. Configure a policy that exports the routes from the routing table into the forwarding table on Router PE3.

```
[edit policy-options]
user@PE3# set policy-statement lb then load-balance per-packet

[edit routing-options]
user@PE3# set forwarding-table export lb
```

Similarly, configure other PE routers.

6. Configure BGP on Router CE2, and include a policy for exporting routes to and from the service provider network.

```
[edit policy-options]
user@CE2# set policy-statement send-direct from protocol direct
user@CE2# set policy-statement send-direct then accept

[edit protocols bgp group toAS2]
user@CE2# set type external
```

```

user@CE2# set export send-direct
user@CE2# set peer-as 64497
user@CE2# set neighbor 10.1.1.25
user@CE2# set neighbor 10.1.1.29

```

Similarly, configure other CE routers.

7. Configure BGP on Router PE3 for routing within the provider core.

```

[edit protocols bgp group toInternal]
user@PE3# set type internal
user@PE3# set family inet-vpn unicast
user@PE3# set family inet6-vpn unicast
user@PE3# set multipath
user@PE3# set local-address 192.0.2.4
user@PE3# set neighbor 192.0.2.2
user@PE3# set neighbor 192.0.2.3

```

Similarly, configure other PE routers.

8. Configure the Layer 3 VPN routing instance on Router PE3.

```

[set routing-instances radium]
user@PE3# set instance-type vrf
user@PE3# set vrf-table-label
user@PE3# set interface ge-2/0/2.0
user@PE3# set interface ge-2/0/3.0
user@PE3# set route-distinguisher 64497:1
user@PE3# set vrf-target target:64497:1

[edit routing-instances radium protocols bgp group toCE2]
user@PE3# set type external
user@PE3# set peer-as 64498
user@PE3# set neighbor 10.1.1.26

[edit routing-instances radium protocols bgp group toCE3]
user@PE3# set type external
user@PE3# set peer-as 64499
user@PE3# set neighbor 10.1.1.22

```

Similarly, configure other PE routers.

9. Configure provider edge link protection on the link between Routers PE3 and CE2.

```

[edit routing-instances radium protocols bgp group toCE2]
user@PE3# set family inet unicast protection
user@PE3# set family inet6 unicast protection

```

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show policy-options**, **show protocols**, and **show routing-instances** commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE3# show interfaces
ge-2/0/0 {
  unit 0 {
    description toPE1;
    family inet {
      address 10.1.1.10/30;
    }
    family inet6 {
      address 2001:db8:0:9::/64 {
        eui-64;
      }
    }
    family mpls;
  }
}

ge-2/0/1 {
  unit 0 {
    description toP;
    family inet {
      address 10.1.1.18/30;
    }
    family inet6 {
      address 2001:db8:0:17::/64 {
        eui-64;
      }
    }
    family mpls;
  }
}

ge-2/0/2 {
  unit 0 {
    description toCE2;
    family inet {
      address 10.1.1.25/30;
    }
    family inet6 {
      address 2001:db8:0:25::/64 {
        eui-64;
      }
    }
    family mpls;
  }
}

ge-2/0/3 {
  unit 0 {
    description toCE3;
    family inet {
      address 10.1.1.21/30;
    }
    family inet6 {
      address 2001:db8:0:21::/64 {
        eui-64;
      }
    }
    family mpls;
  }
}
```

```

    }
  }

  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.4/24;
      }
      family inet6 {
        address 2001:db8::4/128;
      }
    }
  }
}

```

```

user@PE3# show routing-options
router-id 192.0.2.4;
autonomous-system 64497;
forwarding-table {
  export lb;
}

```

```

user@PE3# show policy-options
policy-statement lb {
  then {
    load-balance per-packet;
  }
}

```

```

user@PE3# show protocols
mpls {
  interface all;
}
bgp {
  group toInternal {
    type internal;
    local-address 192.0.2.4;
    family inet-vpn {
      unicast;
    }
    family inet6-vpn {
      unicast;
    }
    multipath;
    neighbor 192.0.2.2;
    neighbor 192.0.2.3;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface ge-2/0/1.0 {
      metric 5;
    }
    interface ge-2/0/0.0 {
      metric 10;
    }
  }
}

```

```

}
ospf3 {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface ge-2/0/1.0 {
      metric 5;
    }
    interface ge-2/0/0.0 {
      metric 10;
    }
  }
}
ldp {
  interface all;
}

user@PE3# show routing-instances
radium {
  instance-type vrf;
  interface ge-2/0/2.0;
  interface ge-2/0/3.0;
  route-distinguisher 64497:1;
  vrf-target target:64497:1;
  protocols {
    bgp {
      group toCE2 {
        type external;
        family inet {
          unicast {
            protection;
          }
        }
        family inet6 {
          unicast {
            protection;
          }
        }
        peer-as 64498;
        neighbor 10.1.1.26;
      }
      group toCE3 {
        type external;
        peer-as 64499;
        neighbor 10.1.1.22;
      }
    }
  }
}

```

Run these commands on all other routers to confirm the configurations. If you are done configuring the routers, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying BGP on page 891](#)
- [Verifying Provider Edge Link Protection on page 892](#)

### Verifying BGP

**Purpose** Verify that BGP is functional in the Layer 3 VPN.

**Action** From operational mode on Router PE3, run the **show route protocol bgp** command.

```
user@PE3> show route protocol bgp
inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

radius.inet.0: 9 destinations, 14 routes (9 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

192.0.2.1/24      *[BGP/170] 00:09:15, localpref 100, from 192.0.2.2
                  AS path: 64496 I, validation-state: unverified
                  > to 10.1.1.9 via ge-2/0/0.0, Push 299792
192.0.2.6/24      @[BGP/170] 00:09:40, localpref 100
                  AS path: 64498 I, validation-state: unverified
                  > to 10.1.1.26 via ge-2/0/2.0
                  [BGP/170] 00:09:07, localpref 100, from 192.0.2.3
                  AS path: 64498 I, validation-state: unverified
                  > to 10.1.1.17 via ge-2/0/1.0, Push 299792, Push 299776(top)
192.0.2.7/24      *[BGP/170] 00:09:26, localpref 100
                  AS path: 64499 I, validation-state: unverified
                  > to 10.1.1.22 via ge-2/0/3.0
10.1.1.0/30       *[BGP/170] 00:09:15, localpref 100, from 192.0.2.2
                  AS path: I, validation-state: unverified
                  > to 10.1.1.9 via ge-2/0/0.0, Push 299792
10.1.1.20/30      [BGP/170] 00:09:26, localpref 100
                  AS path: 64499 I, validation-state: unverified
                  > to 10.1.1.22 via ge-2/0/3.0
10.1.1.24/30      [BGP/170] 00:09:40, localpref 100
                  AS path: 64498 I, validation-state: unverified
                  > to 10.1.1.26 via ge-2/0/2.0
10.1.1.28/30      *[BGP/170] 00:09:07, localpref 100, from 192.0.2.3
                  AS path: I, validation-state: unverified
                  > to 10.1.1.17 via ge-2/0/1.0, Push 299792, Push 299776(top)

                  [BGP/170] 00:09:40, localpref 100
                  AS path: 64498 I, validation-state: unverified
                  > to 10.1.1.26 via ge-2/0/2.0

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

64497:1:192.0.2.1/24
                  *[BGP/170] 00:09:15, localpref 100, from 192.0.2.2
```

```

        AS path: 64496 I, validation-state: unverified
        > to 10.1.1.9 via ge-2/0/0.0, Push 299792
64497:1:192.0.2.6/24
    *[BGP/170] 00:09:07, localpref 100, from 192.0.2.3
        AS path: 64498 I, validation-state: unverified
        > to 10.1.1.17 via ge-2/0/1.0, Push 299792, Push 299776(top)
64497:1:10.1.1.0/30
    *[BGP/170] 00:09:15, localpref 100, from 192.0.2.2
        AS path: I, validation-state: unverified
        > to 10.1.1.9 via ge-2/0/0.0, Push 299792
64497:1:10.1.1.28/30
    *[BGP/170] 00:09:07, localpref 100, from 192.0.2.3
        AS path: I, validation-state: unverified
        > to 10.1.1.17 via ge-2/0/1.0, Push 299792, Push 299776(top)

```

inet6.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

radium.inet6.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)

The output shows all the BGP routes in the routing table of Router PE3. This indicates that BGP is functioning as required.

Similarly, run this command on other routers to check if BGP is operational.

**Meaning** BGP is functional in the Layer 3 VPN.

### *Verifying Provider Edge Link Protection*

**Purpose** Verify that the provider edge link between Routers PE2 and CE2 is protected.

**Action** To verify that provider edge link protection is configured correctly:

1. Confirm that a route on Router CE2 is advertised to Router PE3, directly and through Router PE2.

If the route is advertised correctly, you will see multiple paths for the route.

From operational mode on Router PE3, run the **show route destination-prefix** command.

```

user@PE3> show route 192.0.2.6
radium.inet.0: 9 destinations, 14 routes (9 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

192.0.2.6/24          @[BGP/170] 02:55:36, localpref 100
                     AS path: 64498 I, validation-state: unverified
                     > to 10.1.1.26 via ge-2/0/2.0
                     [BGP/170] 00:10:13, localpref 100, from 192.0.2.3
                     AS path: 64498 I, validation-state: unverified
                     > to 10.1.1.17 via ge-2/0/1.0, Push 299840, Push 299776(top)

                     #[Multipath/255] 00:10:13
                     > to 10.1.1.26 via ge-2/0/2.0
                       to 10.1.1.17 via ge-2/0/1.0, Push 299840, Push 299776(top)

```



The output verifies the presence of multiple paths from Router PE3 to the destination route, **192.0.2.6**, on Router CE2. The first path is directly through the PE3-CE2 link (**10.1.1.26**). The second path is through the provider core and PE2 (**10.1.1.17**).

2. Verify that the protection path is correctly configured by confirming that the weight for the active path being protected is **0x1**, and the weight for the protection candidate path is **0x4000**.

From operational mode on Router PE3, run the **show route destination-prefix extensive** command.

```
user@PE3> show route 192.0.2.6 extensive
radium.inet.0: 9 destinations, 14 routes (9 active, 0 holddown, 0 hidden)
192.0.2.6/24 (3 entries, 2 announced)
    State: <CalcForwarding>
TSI:
KRT in-kernel 192.0.2.6/24 -> {list:10.1.1.26, indirect(1048584)}
Page 0 idx 1 Type 1 val 9229c38
    Nexthop: Self
    AS path: [64497] 64498 I
    Communities:
Page 0 idx 2 Type 1 val 9229cc4
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [64497] 64498 I
    Communities: target:64497:1
Path 192.0.2.6 from 10.1.1.26 Vector len 4. Val: 1 2
    @BGP    Preference: 170/-101
            Next hop type: Router, Next hop index: 994
            Address: 0x9240a74
            Next-hop reference count: 5
            Source: 10.1.1.26
            Next hop: 10.1.1.26 via ge-2/0/2.0, selected
            Session Id: 0x200001
            State: <Active Ext ProtectionPath ProtectionCand>
            Peer AS: 64498
            Age: 2:55:54
            Validation State: unverified
            Task: BGP_64498.10.1.1.26+52214
            Announcement bits (1): 2-BGP_RT_Background
            AS path: 64498 I
            Accepted
            Localpref: 100
            Router ID: 192.0.2.6
    BGP    Preference: 170/-101
            Route Distinguisher: 64497:1
            Next hop type: Indirect
            Address: 0x92413a8
            Next-hop reference count: 6
            Source: 192.0.2.3
            Next hop type: Router, Next hop index: 1322
            Next hop: 10.1.1.17 via ge-2/0/1.0, selected
            Label operation: Push 299840, Push 299776(top)
            Label TTL action: prop-ttl, prop-ttl(top)
            Session Id: 0x200005
            Protocol next hop: 192.0.2.3
            Push 299840
            Indirect next hop: 94100ec 1048584 INH Session ID: 0x20000b
```

```

State: <Secondary NotBest Int Ext ProtectionCand>
Inactive reason: Not Best in its group - Interior > Exterior > Exterior via

Interior
Local AS: 64497 Peer AS: 64497
Age: 10:31 Metric2: 1
Validation State: unverified
Task: BGP_64497.192.0.2.3+179
Local AS: 64497 Peer AS: 64497
Age: 10:31 Metric2: 1
Validation State: unverified
Task: BGP_64497.192.0.2.3+179
AS path: 64498 I
Communities: target:64497:1
Import Accepted
VPN Label: 299840
Localpref: 100
Router ID: 192.0.2.3
Primary Routing Table bgp.l3vpn.0
Indirect next hops: 1
  Protocol next hop: 192.0.2.3 Metric: 1
  Push 299840
  Indirect next hop: 94100ec 1048584 INH Session ID:

0x20000b
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.1.1.17 via ge-2/0/1.0
    Session Id: 0x200005
  192.0.2.3/24 Originating RIB: inet.3
  Metric: 1 Node path count: 1
  Forwarding nexthops: 1
    Nexthop:10.1.1.17 via ge-2/0/1.0
#Multipath Preference: 255
Next hop type: List, Next hop index: 1048585
Address: 0x944c154
Next-hop reference count: 2
Next hop: ELNH Address 0x9240a74 weight 0x1, selected
equal-external-internal-type external
  Next hop type: Router, Next hop index: 994
  Address: 0x9240a74
  Next-hop reference count: 5
  Next hop:10.1.1.26 via ge-2/0/2.0
Next hop: ELNH Address 0x92413a8 weight 0x4000
equal-external-internal-type internal
  Next hop type: Indirect
  Address: 0x92413a8
  Next-hop reference count: 6
  Protocol next hop: 192.0.2.3
  Push 299840
  Indirect next hop: 94100ec 1048584 INH Session ID: 0x20000b

  Next hop type: Router, Next hop index: 1322
  Address: 0x9241310
  Next-hop reference count: 4
  Next hop: 10.1.1.17 via ge-2/0/1.0
  Label operation: Push 299840, Push 299776(top)
  Label TTL action: prop-ttl, prop-ttl(top)
State: <ForwardingOnly Int Ext>
Inactive reason: Forwarding use only
Age: 10:31
Validation State: unverified
Task: RT

```

```
Announcement bits (1): 0-KRT
AS path: 64498 I
```

The output shows that the weight (**0x4000**) assigned to the PE3-CE2 path is greater than the weight (**0x1**) assigned to the PE2-CE2 path. This confirms that the PE2-CE2 path is protected by the PE3-CE2 path.

**Meaning** The provider edge link between Routers PE2 and CE2 is protected.

**See Also** • [Understanding Provider Edge Link Protection in Layer 3 VPNs on page 879](#)

## Example: Configuring Provider Edge Link Protection for BGP Labeled Unicast Paths

This example shows how to configure a labeled unicast protection path that can be used in case of a link failure in a carrier-of-carriers topology.

- [Requirements on page 895](#)
- [Overview on page 895](#)
- [Configuration on page 896](#)
- [Verification on page 907](#)

### Requirements

This example uses the following hardware and software components:

- M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers
- Junos OS Release 13.3 or later

### Overview


This example shows how to configure labeled-unicast link protection in a Layer 3 VPN.

#### Topology

In this example, a carrier-of-carriers topology is set up by configuring two customer edge devices and eight service provider edge devices in five autonomous systems. The CE devices are configured in AS100 and AS101. The PE devices are configured in AS200, AS300, and AS201.

[Figure 82 on page 881](#) shows the topology used in this example.

Figure 83: Labeled Unicast Link Protection in a Layer 3 VPN



The aim of this example is to protect the provider edge link between Routers R4 and R3. Protection is configured on the primary link between R4 and R3 such that the traffic can be routed through the backup link (R11 to R10) when the primary link goes down.



**NOTE:** Protection can also be configured on the secondary link between R11 and R10 so that if that link becomes the primary link and the R4-R3 link becomes secondary, the R11-R10 link will be protected as well.

## 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.



**NOTE:** Protection is added to the configuration only after the initial configuration is committed and BGP has installed the best path in the forwarding table.

```

Router R0    set interfaces ge-2/0/0 unit 0 description toR1
              set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.1/30
              set interfaces ge-2/0/0 unit 0 family iso
              set interfaces ge-2/0/0 unit 0 family mpls
              set interfaces lo0 unit 0 family inet address 192.0.2.1/24
              set interfaces lo0 unit 0 family iso address
                47.0005.80ff.f800.0000.0108.0001.0102.5507.2056.00
              set routing-options router-id 192.0.2.1
              set routing-options autonomous-system 100
              set protocols ospf area 0.0.0.0 interface lo0.0 passive
              set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10

Router R1    set interfaces ge-2/0/0 unit 0 description toR0
              set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.2/30
              set interfaces ge-2/0/0 unit 0 family iso
              set interfaces ge-2/0/0 unit 0 family mpls
              set interfaces ge-2/0/1 unit 0 description toR2
              set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.5/30
              set interfaces ge-2/0/1 unit 0 family iso
              set interfaces ge-2/0/1 unit 0 family mpls
              set interfaces lo0 unit 0 family inet address 192.0.2.2/24
              set interfaces lo0 unit 0 family iso address
                47.0005.80ff.f800.0000.0108.0001.0102.5507.2052.00
              set routing-options router-id 192.0.2.2
              set routing-options autonomous-system 200
              set protocols mpls interface all
              set protocols ldp track-igp-metric
              set protocols ldp interface ge-2/0/1.0
              set protocols ldp interface lo0.0
              set protocols ospf traffic-engineering
              set protocols ospf area 0.0.0.0 interface lo0.0 passive
              set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 10
              set protocols bgp group toR8 local-address 192.0.2.2
              set protocols bgp group toR8 type external
              set protocols bgp group toR8 multihop ttl 10
              set protocols bgp group toR8 family inet-vpn unicast
              set protocols bgp group toR8 neighbor 192.0.2.9 peer-as 201
              set policy-options policy-statement child_vpn_routes from protocol bgp
              set policy-options policy-statement child_vpn_routes then accept
              set routing-instances customer-provider-vpn instance-type vrf
              set routing-instances customer-provider-vpn interface ge-2/0/0.0
              set routing-instances customer-provider-vpn route-distinguisher 192.0.2.4:1
              set routing-instances customer-provider-vpn vrf-target target:200:1
              set routing-instances customer-provider-vpn protocols ospf export child_vpn_routes
              set routing-instances customer-provider-vpn protocols ospf area 0.0.0.0 interface
                ge-2/0/0.0

Router R2    set interfaces ge-2/0/0 unit 0 description toR1
              set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.6/30
              set interfaces ge-2/0/0 unit 0 family iso
              set interfaces ge-2/0/0 unit 0 family mpls
              set interfaces ge-2/0/1 unit 0 description toR3
              set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.9/30
              set interfaces ge-2/0/1 unit 0 family iso

```

```

set interfaces ge-2/0/1 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 description toR10
set interfaces ge-2/0/2 unit 0 family inet address 10.1.0.37/30
set interfaces ge-2/0/2 unit 0 family iso
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.3/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2046.00
set routing-options router-id 192.0.2.3
set routing-options autonomous-system 200
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp interface ge-2/0/0.0
set protocols ldp interface ge-2/0/1.0
set protocols ldp interface ge-2/0/2.0
set protocols ldp interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0 metric 10

```

**Router R3**

```

set interfaces ge-2/0/0 unit 0 description toR2
set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.10/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toR4
set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.13/30
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.4/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2045.00
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 200
set protocols mpls traffic-engineering bgp-igp
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp egress-policy from-bgp
set protocols ldp interface ge-2/0/0.0
set protocols ldp interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf export from-bgp
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols bgp group toR4 type external
set protocols bgp group toR4 import send-local
set protocols bgp group toR4 family inet labeled-unicast
set protocols bgp group toR4 export send-local
set protocols bgp group toR4 neighbor 10.1.0.14 peer-as 300
set policy-options policy-statement from-bgp from protocol bgp
set policy-options policy-statement from-bgp then metric add 100
set policy-options policy-statement from-bgp then accept
set policy-options policy-statement send-local term 2 from metric 100
set policy-options policy-statement send-local term 2 then reject

```

```
set policy-options policy-statement send-local then accept
```

```
Router R4    set interfaces ge-2/0/0 unit 0 description toR3
              set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.14/30
              set interfaces ge-2/0/0 unit 0 family iso
              set interfaces ge-2/0/0 unit 0 family mpls
              set interfaces ge-2/0/1 unit 0 description toR5
              set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.17/30
              set interfaces ge-2/0/1 unit 0 family iso
              set interfaces ge-2/0/1 unit 0 family mpls
              set interfaces lo0 unit 0 family inet address 192.0.2.5/24
              set interfaces lo0 unit 0 family iso address
                47.0005.80ff.f800.0000.0108.0001.0102.5507.2049.00
              set policy-options policy-statement 1b then load-balance per-packet
              set routing-options router-id 192.0.2.5
              set routing-options autonomous-system 300
              set routing-options forwarding-table export 1b
              set protocols mpls interface all
              set protocols ldp track-igp-metric
              set protocols ldp interface ge-2/0/1.0
              set protocols ldp interface lo0.0
              set protocols isis level 1 disable
              set protocols isis level 2 wide-metrics-only
              set protocols isis interface ge-2/0/1.0 level 2 metric 10
              set protocols isis interface lo0.0 passive
              set protocols bgp group parent-vpn-peers type internal
              set protocols bgp group parent-vpn-peers local-address 192.0.2.5
              set protocols bgp group parent-vpn-peers family inet-vpn unicast
              set protocols bgp group parent-vpn-peers neighbor 192.0.2.7
              set protocols bgp group parent-vpn-peers neighbor 192.0.2.12
              set routing-instances coc-provider-vpn instance-type vrf
              set routing-instances coc-provider-vpn interface ge-2/0/0.0
              set routing-instances coc-provider-vpn interface ge-2/0/2.0
              set routing-instances coc-provider-vpn route-distinguisher 192.0.2.5:1
              set routing-instances coc-provider-vpn vrf-target target:300:1
              set routing-instances coc-provider-vpn protocols bgp group toR3 type external
              set routing-instances coc-provider-vpn protocols bgp group toR3 family inet
                labeled-unicast per-prefix-label
              set routing-instances coc-provider-vpn protocols bgp group toR3 neighbor 10.1.0.13
                peer-as 200

Router R5    set interfaces ge-2/0/0 unit 0 description toR4
              set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.18/30
              set interfaces ge-2/0/0 unit 0 family iso
              set interfaces ge-2/0/0 unit 0 family mpls
              set interfaces ge-2/0/1 unit 0 description toR6
              set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.21/30
              set interfaces ge-2/0/1 unit 0 family iso
              set interfaces ge-2/0/1 unit 0 family mpls
              set interfaces ge-2/0/2 unit 0 description toR11
              set interfaces ge-2/0/2 unit 0 family inet address 10.1.0.46/30
              set interfaces ge-2/0/2 unit 0 family iso
              set interfaces ge-2/0/2 unit 0 family mpls
              set interfaces lo0 unit 0 family inet address 192.0.2.6/24
```

```

set interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.5507.2050.00
set routing-options router-id 192.0.2.6
set routing-options autonomous-system 300
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp interface ge-2/0/0.0
set protocols ldp interface ge-2/0/1.0
set protocols ldp interface ge-2/0/2.0
set protocols ldp interface lo0.0
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface ge-2/0/0.0 level 2 metric 10
set protocols isis interface ge-2/0/1.0 level 2 metric 10
set protocols isis interface ge-2/0/2.0 level 2 metric 10
set protocols isis interface lo0.0 passive

```

```

Router R6  set interfaces ge-2/0/0 unit 0 description toR5
set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.22/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toR7
set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.25/30
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.7/24
set interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.5507.2048.00
set routing-options router-id 192.0.2.7
set routing-options autonomous-system 300
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp interface ge-2/0/0.0
set protocols ldp interface lo0.0
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface ge-2/0/0.0 level 2 metric 10
set protocols isis interface lo0.0 passive
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.7
set protocols bgp group parent-vpn-peers family inet-vpn unicast
set protocols bgp group parent-vpn-peers neighbor 192.0.2.5
set protocols bgp group parent-vpn-peers neighbor 192.0.2.12
set routing-instances coc-provider-vpn instance-type vrf
set routing-instances coc-provider-vpn interface ge-2/0/1.0
set routing-instances coc-provider-vpn route-distinguisher 192.0.2.7:1
set routing-instances coc-provider-vpn vrf-target target:300:1
set routing-instances coc-provider-vpn protocols bgp group toR7 family inet
  labeled-unicast per-prefix-label
set routing-instances coc-provider-vpn protocols bgp group toR7 type external
set routing-instances coc-provider-vpn protocols bgp group toR7 neighbor 10.1.0.26
  peer-as 201

```

```

Router R7  set interfaces ge-2/0/0 unit 0 description toR6

```



```

set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.26/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toR8
set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.29/30
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.8/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2054.00
set routing-options router-id 192.0.2.8
set routing-options autonomous-system 201
set protocols mpls traffic-engineering bgp-igp
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp egress-policy from-bgp
set protocols ldp interface ge-2/0/1.0
set protocols ldp interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf export from-bgp
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0 metric 10
set protocols bgp group toR6 type external
set protocols bgp group toR6 import send-all
set protocols bgp group toR6 family inet labeled-unicast
set protocols bgp group toR6 export send-all
set protocols bgp group toR6 neighbor 10.1.0.25 peer-as 300
set policy-options policy-statement from-bgp from protocol bgp
set policy-options policy-statement from-bgp then accept
set policy-options policy-statement send-all then accept

```

**Router R8**

```

set interfaces ge-2/0/0 unit 0 description toR7
set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.30/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toR9
set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.33/30
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.9/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2053.00
set routing-options router-id 192.0.2.9
set routing-options autonomous-system 201
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp interface ge-2/0/0.0
set protocols ldp interface lo0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols bgp group toR1 local-address 192.0.2.9
set protocols bgp group toR1 type external
set protocols bgp group toR1 multihop ttl 10
set protocols bgp group toR1 family inet-vpn unicast
set protocols bgp group toR1 neighbor 192.0.2.2 peer-as 200

```

```

set policy-options policy-statement child_vpn_routes from protocol bgp
set policy-options policy-statement child_vpn_routes then accept
set routing-instances customer-provider-vpn instance-type vrf
set routing-instances customer-provider-vpn interface ge-2/0/1.0
set routing-instances customer-provider-vpn route-distinguisher 192.0.2.9:1
set routing-instances customer-provider-vpn vrf-target target:200:1
set routing-instances customer-provider-vpn protocols ospf export child_vpn_routes
set routing-instances customer-provider-vpn protocols ospf area 0.0.0.0 interface
  ge-2/0/1.0

```

**Router R9**

```

set interfaces ge-2/0/0 unit 0 description toR8
set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.34/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.10/24
set interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.5507.2047.00
set routing-options router-id 192.0.2.10
set routing-options autonomous-system 101
set routing-options static route 198.51.100.1/24 discard
set protocols ospf export statics
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set policy-options policy-statement statics from route-filter 198.51.100.1/24 exact
set policy-options policy-statement statics then accept

```

**Router R10**

```

set interfaces ge-2/0/0 unit 0 description toR2
set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.38/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toR11
set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.41/30
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.11.24
set interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.5507.2061.00
set routing-options router-id 192.0.2.11
set routing-options autonomous-system 200
set protocols mpls traffic-engineering bgp-igp
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp egress-policy from-bgp
set protocols ldp interface ge-2/0/0.0
set protocols ldp interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf export from-bgp
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 10
set protocols bgp group toR4 type external
set protocols bgp group toR4 import send-local
set protocols bgp group toR4 family inet labeled-unicast
set protocols bgp group toR4 export send-local
set protocols bgp group toR4 neighbor 10.1.0.42 peer-as 300

```

```

set protocols bgp group toR4 inactive: neighbor 10.1.0.50 peer-as 300
set policy-options policy-statement from-bgp from protocol bgp
set policy-options policy-statement from-bgp then metric add 100
set policy-options policy-statement from-bgp then accept
set policy-options policy-statement send-local term 2 from metric 100
set policy-options policy-statement send-local term 2 then reject
set policy-options policy-statement send-local then accept

```

```

Router R11
set interfaces ge-2/0/0 unit 0 description toR10
set interfaces ge-2/0/0 unit 0 family inet address 10.1.0.42/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/1 unit 0 description toR5
set interfaces ge-2/0/1 unit 0 family inet address 10.1.0.45/30
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.12/24
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5507.2062.00
set routing-options router-id 192.0.2.12
set routing-options autonomous-system 300
set protocols mpls interface all
set protocols ldp track-igp-metric
set protocols ldp interface ge-2/0/1.0
set protocols ldp interface lo0.0
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface ge-2/0/1.0 level 2 metric 10
set protocols isis interface lo0.0 passive
set protocols bgp group parent-vpn-peers type internal
set protocols bgp group parent-vpn-peers local-address 192.0.2.12
set protocols bgp group parent-vpn-peers family inet-vpn unicast
set protocols bgp group parent-vpn-peers neighbor 192.0.2.7
set protocols bgp group parent-vpn-peers neighbor 192.0.2.12
set routing-instances coc-provider-vpn instance-type vrf
set routing-instances coc-provider-vpn interface ge-2/0/0.0
set routing-instances coc-provider-vpn route-distinguisher 192.0.2.12:1
set routing-instances coc-provider-vpn vrf-target target:300:1
set routing-instances coc-provider-vpn protocols bgp group toR10 family inet
    labeled-unicast per-prefix-label
set routing-instances coc-provider-vpn protocols bgp group toR10 type external
set routing-instances coc-provider-vpn protocols bgp group toR10 neighbor 10.1.0.41
    peer-as 200

```

### Configuring Provider Edge Link Protection in Layer 3 VPNs

**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 labeled unicast link protection:

1. Configure the router interfaces.

[edit interfaces]

```
user@R4# set ge-2/0/0 unit 0 description toR3
user@R4# set ge-2/0/0 unit 0 family inet address 10.1.0.14/30
user@R4# set ge-2/0/0 unit 0 family iso
user@R4# set ge-2/0/0 unit 0 family mpls

user@R4# set ge-2/0/1 unit 0 description toR5
user@R4# set ge-2/0/1 unit 0 family inet address 10.1.0.17/30
user@R4# set ge-2/0/1 unit 0 family iso
user@R4# set ge-2/0/1 unit 0 family mpls

user@R4# set lo0 unit 0 family inet address 192.0.2.5/24
user@R4# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5507.2049.00
```

Similarly, configure the interfaces on all other routers.

2. Configure the routing policy options on R4.

```
[edit policy-options]
user@R4# set policy-statement 1b then load-balance per-packet
```

Similarly, configure the policy options on routers R1, R3, R7, R8, R9, and R10 for this example.

3. Configure the router ID, autonomous system (AS) number, and any other routing options.

```
[edit routing-options]
user@R4# set router-id 192.0.2.5
user@R4# set autonomous-system 300
user@R4# set forwarding-table export 1b
```

Similarly, configure the router ID, AS number, and any other routing options for all other routers. In this example, the router ID is chosen to be identical to the loopback address configured on the router.

4. Configure MPLS and LDP on Router R4.

```
[edit protocols]
user@R4# set mpls interface all
user@R4# set ldp track-igp-metric
user@R4# set ldp interface ge-2/0/1.0
user@R4# set ldp interface lo0.0
```

Similarly, configure MPLS and LDP on all other routers except R0 and R9.

5. Configure an IGP on the core-facing interfaces of Router R4.

```
[edit protocols isis]
user@R4# set level 1 disable
user@R4# set level 2 wide-metrics-only
user@R4# set interface ge-2/0/1.0 level 2 metric 10
user@R4# set interface lo0.0 passive
```

Similarly, configure other routers (IS-IS on R5, R6, and R11 and OSPF on all other routers in this example).

6. Configure BGP on Router R4.

```
[edit protocols bgp group parent-vpn-peers]
user@R4# set type internal
user@R4# set local-address 192.0.2.5
user@R4# set family inet-vpn unicast
user@R4# set neighbor 192.0.2.7
user@R4# set neighbor 192.0.2.12
```

Similarly, configure BGP on routers R1, R3, R6, R7, R8, R10, and R11.

7. Configure a VPN routing and forwarding (VRF) instance on Router R4 to create a Layer 3 VPN.

```
[edit routing-instances coc-provider-vpn]
user@R4# set instance-type vrf
user@R4# set interface ge-2/0/0.0
user@R4# set interface ge-2/0/2.0
user@R4# set route-distinguisher 192.0.2.5:1
user@R4# set vrf-target target:300:1

[edit routing-instances coc-provider-vpn protocols bgp group toR3]
user@R4# set type external
user@R4# set family inet labeled-unicast per-prefix-label
user@R4# set neighbor 10.1.0.13 peer-as 200
```

Similarly, configure other VRF routing instances on R1, R6, R8, and R11.

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show routing-options**, **show protocols**, and **show routing-instances** commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
ge-2/0/0 {
  unit 0 {
    description toR3;
    family inet {
      address 10.1.0.14/30;
    }
    family iso;
    family mpls;
  }
}
ge-2/0/1 {
  unit 0 {
    description toR5;
    family inet {
      address 10.1.0.17/30;
    }
    family iso;
    family mpls;
  }
}
```

```
    }  
  }  
  
  lo0 {  
    unit 0 {  
      family inet {  
        address 192.0.2.5/24;  
      }  
      family iso {  
        address 47.0005.80ff.f800.0000.0108.0001.0102.5507.2049.00;  
      }  
    }  
  }  
}  
  
user@R4# show policy-options  
policy-statement lb {  
  then {  
    load-balance per-packet;  
  }  
}  
  
user@R4# show routing-options  
router-id 192.0.2.5;  
autonomous-system 300;  
forwarding-table {  
  export lb;  
}  
  
user@R4# show protocols  
mpls {  
  interface all;  
}  
ldp {  
  track-igp-metric;  
  interface ge-2/0/1.0;  
  interface lo0.0;  
}  
isis {  
  level 1 disable;  
  level 2 wide-metrics-only;  
  interface ge-2/0/1.0 {  
    level 2 metric 10;  
  }  
  interface lo0.0 {  
    passive;  
  }  
}  
bgp {  
  group parent-vpn-peers {  
    type internal;  
    local-address 192.0.2.5;  
    family inet-vpn {  
      unicast;  
    }  
    neighbor 192.0.2.7;  
    neighbor 192.0.2.12;  
  }  
}
```

```

user@R4# show routing-instances
coc-provider-vpn {
  instance-type vrf;
  interface ge-2/0/0.0;
  interface ge-2/0/2.0;
  route-distinguisher 192.0.2.5:1;
  vrf-target target:300:1;
  protocols {
    bgp {
      group toR3 {
        type external;
        family inet {
          labeled-unicast {
            per-prefix-label;
          }
        }
        neighbor 10.1.0.13 {
          peer-as 200;
        }
      }
    }
  }
}

```

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

Repeat the procedure for every router in this example, using the appropriate interface names and addresses for each router.

## Verification

Confirm that the configuration is working properly.

- [Enabling Protection on page 907](#)
- [Verifying Multipath Entries on page 908](#)
- [Verifying That Multipath Entries Have Different Weights on page 908](#)

### Enabling Protection

**Purpose** Enable protection on R4 to request protection for the link from R4 to R3.

**Action** 1. Add the **protection** statement at the **[edit routing-instances *instance-name* protocols bgp group *group-name* family inet labeled-unicast]** hierarchy level.

```

[edit routing-instances coc-provider-vpn protocols bgp group toR3]
user@R4# set family inet labeled-unicast protection

```

2. Verify and commit the configuration.

```

type external;
family inet {
  labeled-unicast {
    per-prefix-label;
  }
}

```

```

        protection;
    }
}
neighbor 10.1.0.13 {
    peer-as 200;
}

```

### Verifying Multipath Entries

**Purpose** Verify that R4 has a multipath entry with two entries.

**Action** From operational mode on Router R4, run the **show route 192.0.2.2** command to check the route to R1.

```

user@R4> show route 192.0.2.2
#[Multipath/255] 00:02:44, metric 20
  > to 10.1.0.13 via ge-2/0/0.0, Push 408592
    to 10.1.0.18 via ge-2/0/1.0, Push 299856, Push 299792(top)

```

### Verifying That Multipath Entries Have Different Weights

**Purpose** Verify that the two routes in the multipath entry have different weights, with the first entry having a weight of 0x1 and the second having a weight of 0x4000.

**Action** From operational mode on Router R4, run the **show route 192.0.2.2 detail** command to check the route to R1.

```

user@R4> show route 192.0.2.2 detail
#Multipath Preference: 255
  Next hop type: List, Next hop index: 1048609
  Address: 0x92f058c
  Next-hop reference count: 4
  Next hop: ELNH Address 0x92c48ac weight 0x1, selected
    equal-external-internal-type external
      Next hop type: Router, Next hop index: 1603
      Address: 0x92c48ac
      Next-hop reference count: 2
      Next hop: 10.1.0.13 via ge-2/0/0.0
      Label operation: Push 408592
      Label TTL action: prop-ttl
  Next hop: ELNH Address 0x92c548c weight 0x4000
    equal-external-internal-type internal
      Next hop type: Indirect
      Address: 0x92c548c
      Next-hop reference count: 3
      Protocol next hop: 192.0.2.12
      Push 299856
      Indirect next hop: 0x9380f40 1048608 INH Session ID: 0x10001a

      Next hop type: Router, Next hop index: 1586
      Address: 0x92c5440
      Next-hop reference count: 3
      Next hop: 10.1.0.18 via ge-2/0/1.0

```



```
Label operation: Push 299856, Push 299792(top)
Label TTL action: prop-ttl, prop-ttl(top)
State: <ForwardingOnly Int Ext>
Inactive reason: Forwarding use only
Age: 3:38      Metric: 20
Validation State: unverified
Task: RT
Announcement bits (1): 0-KRT
AS path: 200 I
```

**Related Documentation** • [Understanding Provider Edge Link Protection for BGP Labeled Unicast Paths on page 877](#)



# Improving Layer 3 VPN Performance

- [Example: Configuring Link Protection with Host Fast Reroute on page 911](#)
- [Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928](#)

## Example: Configuring Link Protection with Host Fast Reroute

---

- [Understanding Host Fast Reroute on page 911](#)
- [Example: Configuring Link Protection with Host Fast Reroute on page 916](#)


### Understanding Host Fast Reroute

Host fast reroute (HFRR) adds a precomputed protection path into the Packet Forwarding Engine (PFE), such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the PFE can use another path without having to wait for the router or the protocols to provide updated forwarding information. This precomputed protection path is often called a repair or a backup path.

HFRR is a technology that protects IP endpoints on multipoint interfaces, such as Ethernet. This technology is important in datacenters where fast service restoration for server endpoints is critical. After an interface or a link goes down, HFRR enables the local repair time to be approximately 50 milliseconds.

Consider the network topology shown in [Figure 84 on page 912](#).

Figure 84: Host Fast Reroute



g041181

Routing devices create host route forwarding entries triggered by the Address Resolution Protocol (ARP) and IPv6 Neighbor Discovery Protocol (NDP). HFRR augments the host routes with backup next hops supplied by routing protocols. These backup next hops enable arriving traffic to keep flowing while the network reconverges.

Traffic flows from networks connected to the provider edge devices, PE1 and PE2, to host A and host B. This traffic is protected with HFRR. If the link goes down between device PE2 and the host servers, traffic is rerouted through device PE1 to the host servers. In the topology, host A and host B represent LAN PCs, collectively known as a server farm. The PE devices are routers with a Layer 3 VPN configured between them. Device PE1 learns about the directly connected hosts by way of ARP or the IPv6 NDP.

Device PE2 also has information about the server farm network and advertises this information to Device PE1. This advertisement is transmitted through the Layer 3 VPN using internal BGP (IBGP). On Devices PE1 and PE2, this route is considered a direct route to the server farm subnet.

Device PE1 uses the host routes learned through ARP and NDP to send traffic to the host machines in the server farm. If the link between Device PE1 and the server farm is disrupted and if HFRR is not configured, the routing device finds the next best route, which is the IBGP route. This implementation results in traffic loss for an interval until the update occurs and the network reconverges. HFRR configured on Device PE1 resolves this issue by augmenting the ARP and NDP routes with a backup path so that traffic can continue to be forwarded without interruption.

The backup path in this particular topology is the IBGP Layer 3 VPN route. In an actual deployment, Device PE2 can also configure link protection for its directly connected server farm network, and Device PE1 can advertise reachability to the server farm through itself using the Layer 3 VPN routes to Device PE2. Therefore, HFRR should be enabled on both Device PE1 and Device PE2. Also, Device PE1 and Device PE2 should both advertise reachability to the server farm through BGP.

A temporary routing loop can develop between the PE devices if, for example, the link between Device PE1 to the server farm and the link between Device PE2 to the server farm both go down at same time. The loop can continue until BGP on both ends learns that the server farm subnet is down and withdraws the BGP routes.

### ARP Prefix Limit and Blackout Supplementary Timeout

When you configure HFRR profiles, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore, FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices. The ARP prefix limit does not limit ARP routes in the forwarding table. It does, however, limit the number of ARP routes that Junos OS reads for a profile and therefore limits the number of HFRR routes that the routing process (rpd) creates in the routing table and the forwarding table.

The ARP prefix limit is applied to each HFRR profile. It does not limit the total count of all ARP/HFRR routes in the routing table. It only limits the number of ARP/HFRR routes for each HFRR profile.

There are two configuration statements (`global-arp-prefix-limit` and `arp-prefix-limit`) that set the ARP prefix limit, one at the global `[edit routing-options host-fast-reroute]` hierarchy level and the other at the `[edit routing-instances instance-name routing-options interface interface-name]` hierarchy level, respectively. The global `global-arp-prefix-limit` statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The `arp-prefix-limit` statement overrides the `global-arp-prefix-limit` for that HFRR profile for that protected interface.

When the number of ARP routes in an HFRR profile reaches 80% of the configured ARP prefix limit, a warning message is sent to the system log. The warning message is displayed for any subsequent ARP route added to that HFRR profile if the ARP prefixes remain at greater than 80% of the configured value.

When the number of ARP routes in an HFRR profile reaches 100% of the configured ARP prefix limit for an HFRR profile, another warning message is sent to the system log. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.

After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.

There are global and per-HFRR CLI statements (`global-supplementary-blackout-timer` and `supplementary-blackout-timer`). The global value is at the `[edit routing-options host-fast-reroute]` hierarchy level and applies to all HFRR profiles on the routing device. The supplementary blackout timer configured for the routing-instance interface at the `[edit routing-instances instance-name routing-options interface interface-name]` hierarchy level overrides the global value for that HFRR profile only.

When the blackout timer expires, the HFRR profile is reactivated, and Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.

If an HFRR profile is blacklisted and is in the deactivated state, a reevaluation of the ARP state is performed during every commit operation or whenever the routing process (rpd) is restarted with the **restart routing** command.

### Primary Route and Backup Route Candidates

---

The primary route for the HFRR next hop is provided by the ARP and IPv6 NDP routes. These are /32 or /128 routes. The backup route is an exact prefix match of the address configured on the local interface. For example, if the local address configured is 10.0.0.5/24, the routing device looks for an exact match of prefix 10.0.0.0 with a prefix length of 24 for selection of backup route.

Constraints for backup route selection are as follows:

- Must be a prefix matching the same subnet address configured on the routing device's HFRR-enabled interface.
- The remote end must not have route aggregation (also known as summarization) configured. For example, if the remote end combines two or more /24 subnets to advertise a subnet with a prefix length smaller than /24, the Junos OS does not select this summarized route to be a backup route.
- If there is another route in the routing table learned by another protocol with a longest-prefix match for the /32 or /128 (ARP or NDP) route, that route is not selected to be a backup candidate. For example, suppose that the local interface address is 10.0.0.5/24. Also suppose that the routing table contains an IBGP route with a prefix of 10.0.0.0/24 and an OSPF route with a prefix of 10.0.0.0/28. Even though the /28 route is a better route for certain prefixes in the subnet, the Junos OS does not consider 10.0.0.0/28 to be a backup candidate. The IBGP route becomes the backup candidate for all host routes. However, after the global repair, the OSPF route is used for forwarding.

In short, the backup candidate must be a route with the same prefix as the subnet local interface that you are protecting with HFRR.

### Backup Path Selection Policy

---

Only Layer 3 VPN routes are considered for backup selection. HFRR uses the usual BGP path selection algorithm to select one best backup route. Only one backup path is selected. In case there are multiple backup path candidates, the selection algorithm selects the best backup path. HFRR provides only two paths, one primary and one backup at any point in time. If the selected backup path itself has two paths in it, then the first path in that backup next hop is used as the backup next hop for the HFRR route.

The primary path is installed with a weight of one. The backup path is installed with a weight of 0x4000. The backup path obviously must be a path through an interface that is not the same as the primary interface.

The backup route is looked up only in the routing table to which interface belongs. For IPv4, the Junos OS uses *routing-instance-name.inet.0*. For IPv6, the Junos OS looks in *routing-instance-name.inet6.0*.

---

### Characteristics of HFRR Routes

---

The HFRR route is a forwarding-only route and is not used for route resolution. HFRR routes have host addresses, meaning that they have /32 or /128 as the prefix length. In the case of platforms with dual Routing Engines, the backup routing process (rpd) also creates HFRR routes. However, the backup routing process (rpd) does not install HFRR routes to a routing table until the backup becomes the master after a Routing Engine switchover.

Also note that if an HFRR route is present in the routing table, the HFRR route is used for the unicast reverse-path-forwarding (uRPF) computation.

---

### Removal of HFRR Routes

---

HFRR routes are deleted if the protected interface is deleted or deactivated in the configuration, if HFRR is configured on a routing instance and the routing instance is deactivated or deleted, or when the statement that enables HFRR ([link-protection \(Host Fast Reroute\)](#)) is deleted or deactivated. HFRR routes are deleted and readded when there is a catastrophic operation on routing the instance, such as when the routing process is restarted. HFRR routes are also be removed if all backup routes are deleted. such as when BGP withdraws routes or when BGP is deactivated or deleted.

After a protected interface goes down and if HFRR is deleted or deactivated, a timer starts with a timeout of 20 seconds. The HFRR route deletion occurs after the timer expires. This is to ensure that if the interface is flapping (quickly going up and down), the Junos OS does not unnecessarily perform route deletions and additions that cause traffic loss. This timer is used only when the interface is down or when the HFRR route is deleted or deactivated.

HFRR routes are purged immediately in the following cases:

- A backup route goes down and there are no other potential backup paths.
- An ARP delete message is received.
- The routing process (rpd) terminates.

---

### Interfaces That Support HFRR

---

HFRR is allowed only on Ethernet interfaces. The commit operation fails if you configure HFRR on point-to-point interfaces.

Only interfaces configured under routing instance of type VPN routing and forwarding (VRF) are accepted. The commit operation fails if you configure HFRR on other types of routing instances.

When the following requirements are not met, the commit operation does not fail. However, the interface is not protected by HFRR, and the interface is marked inactive in the [show hfr profiles](#) command output:

- HFRR is allowed only on numbered interfaces, meaning that an address must be assigned to the interface. You cannot, for example, configure IPv4 on the interface with an address and IPv6 without an address.

- Interfaces that are configured for HFRR protection must be configured at the **[edit interfaces]** hierarchy level and also must be attached to the routing instance.
- The routing instance must have a virtual tunnel (VT) interface or the **vrf-table-label** statement included.

Another reason the interface might be marked inactive in the **show hfr profiles** command output is when the interface is migrating from one instance to another, and the HFRR configuration is in the previous routing instance.

HFRR is not supported on overlapping logical units if they belong to the same routing instance, as shown here:

```
user@host # show interfaces
ge-0/0/2 {
  vlan-tagging;
  unit 0 {
    vlan-id 1;
    family inet {
      address 172.16.0.4/16; # same subnet
    }
  }
  unit 1 {
    vlan-id 2;
    family inet {
      address 172.16.0.5/16; # same subnet
    }
  }
}
```

If you configure overlapping subnets as shown here, and if you enable HFRR on both of the overlapping subnets, the routing protocol process (rpd) generates an RPD\_ASSERT error.

- See Also**
- [Example: Configuring Link Protection with Host Fast Reroute on page 916](#)
  - *Understanding BGP Path Selection*

## Example: Configuring Link Protection with Host Fast Reroute

This example shows you how to configure host fast reroute (HFRR). HFRR protects IP endpoints on multipoint interfaces, such as Ethernet.

- [Requirements on page 917](#)
- [Overview on page 917](#)
- [Configuration on page 918](#)
- [Verification on page 925](#)



---

## Requirements

This example uses the following hardware and software components:

- Two provider edge (PE) devices and four provider (P) devices.
- The example assumes that hosts are present, behind the PE devices.
- The example assumes that at least one Layer 3 switch, such as an EX Series switch, is attached to the hosts.
- Junos OS 11.4R2 or later.

---

## Overview

In this example, traffic flows to server hosts from networks connected to PE devices. This traffic is protected with HFRR. If the link goes down between one PE device and the server farm, traffic is rerouted to the server farm through the other PE device.

You can configure HFRR by adding the **link-protection** statement to the interface configuration in the routing instance.

```
[edit routing-instances cust1 routing-options]  
set interface ge-4/1/0.0 link-protection (Host Fast Reroute)
```

We recommend that you include this statement on all PE devices that are connected to server farms through multipoint interfaces.


In this example, the PE devices advertise reachability to their server farms through Layer 3 VPN routes and BGP.

As optional settings, the PE devices have the high availability features Nonstop Active Routing and Virtual Router Redundancy Protocol (VRRP) configured. Nonstop active routing (NSR) enables a routing platform with redundant Routing Engines to switch from a primary Routing Engine to a backup Routing Engine without alerting peer nodes that a change has occurred and without losing routing and protocol information. VRRP provides for automatic assignment of available routers to participating hosts, thus increasing the availability and reliability of routing paths.

## Topology Diagram

Figure 85 on page 918 shows the topology used in this example.

Figure 85: Host Fast Reroute Topology



g041180

This example shows the configuration on all of the routing devices and shows the step-by-step procedure on Device PE1.

### 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 PE1
set interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-4/1/0 unit 0 description toPE2
set interfaces ge-0/2/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-0/2/0 unit 0 description toP1
set interfaces ge-0/2/0 unit 0 family mpls
set interfaces ge-0/2/4 unit 0 family inet address 10.10.15.2/30
set interfaces ge-0/2/4 unit 0 description toP5
set interfaces ge-0/2/4 unit 0 family mpls
set interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24 vrrp-group 1 virtual-address 192.0.2.5
set interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24 vrrp-group 1 priority 240
set interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24 vrrp-group 1 fast-interval 100
set interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24 vrrp-group 1 preempt
```

```

set interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24 vrrp-group 1 accept-data
set interfaces lo0 unit 0 family inet address 10.255.8.207/32
set protocols mpls interface ge-0/2/0.0
set protocols mpls interface ge-0/2/4.0
set protocols bgp group pe-ce type internal
set protocols bgp group pe-ce local-address 10.255.8.207
set protocols bgp group pe-ce family inet-vpn unicast
set protocols bgp group pe-ce neighbor 10.255.8.86
set protocols bgp group pe-ce export send-routes
set protocols ospf area 0.0.0.0 interface ge-0/2/0.0
set protocols ospf area 0.0.0.0 interface ge-0/2/4.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/2/0.0
set protocols ldp interface ge-0/2/4.0
set policy-options policy-statement send-routes term 1 from protocol direct
set policy-options policy-statement send-routes term 1 from protocol local
set policy-options policy-statement send-routes term 1 then accept
set routing-options nonstop-routing
set routing-options autonomous-system 100
set routing-instances cust1 instance-type vrf
set routing-instances cust1 interface ge-4/1/0.0
set routing-instances cust1 route-distinguisher 100:100
set routing-instances cust1 vrf-target target:100:100
set routing-instances cust1 vrf-table-label
set routing-instances cust1 routing-options interface ge-4/1/0.0 link-protection

```

**Device PE2**

```

set interfaces ge-0/0/2 unit 0 family inet address 10.10.12.2/30
set interfaces ge-0/0/2 unit 0 description toP2
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/1/2 unit 0 family inet address 10.10.13.1/30
set interfaces ge-0/1/2 unit 0 description toP4
set interfaces ge-0/1/2 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.3/24
set interfaces ge-2/0/2 unit 0 description toPE1
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.3/24 vrrp-group 1 virtual-address
192.0.2.5
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.3/24 vrrp-group 1 fast-interval
100
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.3/24 vrrp-group 1 preempt
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.3/24 vrrp-group 1 accept-data
set interfaces lo0 unit 0 family inet address 10.255.8.86/32
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/1/2.0
set protocols bgp group pe-ce type internal
set protocols bgp group pe-ce export send-routes
set protocols bgp group pe-ce local-address 10.255.8.86
set protocols bgp group pe-ce family inet-vpn unicast
set protocols bgp group pe-ce neighbor 10.255.8.207
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/1/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface ge-0/1/2.0
set policy-options policy-statement send-routes term 1 from protocol direct
set policy-options policy-statement send-routes term 1 from protocol local

```

```

set policy-options policy-statement send-routes term 1 then accept
set routing-options nonstop-routing
set routing-options autonomous-system 100
set routing-instances cust1 instance-type vrf
set routing-instances cust1 interface ge-2/0/2.0
set routing-instances cust1 route-distinguisher 100:100
set routing-instances cust1 vrf-target target:100:100
set routing-instances cust1 vrf-table-label
set routing-instances cust1 routing-options interface ge-2/0/2.0 link-protection

```

**Device P1**

```

set interfaces ge-0/0/3 unit 0 family inet address 10.10.11.1/30
set interfaces ge-0/0/3 unit 0 description toP2
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family inet address 10.10.10.2/30
set interfaces ge-0/0/4 unit 0 description toPE1
set interfaces ge-0/0/4 unit 0 family mpls
set protocols mpls interface ge-0/0/4.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/0/3.0
set protocols ldp interface ge-0/0/4.0
set routing-options autonomous-system 100

```

**Device P2**

```

set interfaces ge-0/2/1 unit 0 family inet address 10.10.12.1/30
set interfaces ge-0/2/1 unit 0 description toPE2
set interfaces ge-0/2/1 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.10.11.2/30
set interfaces ge-1/2/1 unit 0 description toP1
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.8.246/32
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-options autonomous-system 100

```

**Device P4**

```

set interfaces ge-0/2/3 unit 0 family inet address 10.10.13.2/30
set interfaces ge-0/2/3 unit 0 description toPE2
set interfaces ge-0/2/3 unit 0 family mpls
set interfaces ge-1/3/3 unit 0 family inet address 10.10.14.1/30
set interfaces ge-1/3/3 unit 0 description toP5
set interfaces ge-1/3/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.8.4/32
set protocols mpls interface ge-0/2/3.0
set protocols mpls interface ge-1/3/3.0
set protocols ospf area 0.0.0.0 interface ge-0/2/3.0
set protocols ospf area 0.0.0.0 interface ge-1/3/3.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/2/3.0
set protocols ldp interface ge-1/3/3.0

```

```
set routing-options autonomous-system 100
```

**Device P5**

```
set interfaces ge-0/1/2 unit 0 family inet address 10.10.15.1/30
set interfaces ge-0/1/2 unit 0 description toPE1
set interfaces ge-0/1/2 unit 0 family mpls
set interfaces ge-0/1/5 unit 0 family inet address 10.10.14.2/30
set interfaces ge-0/1/5 unit 0 description toP4
set interfaces ge-0/1/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.8.5/32
set protocols mpls interface ge-0/1/5.0
set protocols mpls interface ge-0/1/2.0
set protocols ospf area 0.0.0.0 interface ge-0/1/5.0
set protocols ospf area 0.0.0.0 interface ge-0/1/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/1/2.0
set protocols ldp interface ge-0/1/5.0
set routing-options autonomous-system 100
```

**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 HFRR:

1. Configure the interfaces.

```
[edit interfaces]
user@PE1# set ge-4/1/0 unit 0 family inet address 192.0.2.2/24
user@PE1# set ge-4/1/0 unit 0 description toPE2
user@PE1# set ge-0/2/0 unit 0 family inet address 10.10.10.1/30
user@PE1# set ge-0/2/0 unit 0 description toP1
user@PE1# set ge-0/2/0 unit 0 family mpls
user@PE1# set ge-0/2/4 unit 0 family inet address 10.10.15.2/30
user@PE1# set ge-0/2/4 unit 0 description toP5
user@PE1# set ge-0/2/4 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 10.255.8.207/32
```

2. (Optional) Configure VRRP on the interface to Device PE2.

```
[edit interfaces ge-4/1/0 unit 0 family inet address 192.0.2.2/24]
user@PE1# set vrrp-group 1 virtual-address 192.0.2.5
user@PE1# set vrrp-group 1 priority 240
user@PE1# set vrrp-group 1 fast-interval 100
user@PE1# set vrrp-group 1 preempt
user@PE1# set vrrp-group 1 accept-data
```

3. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PE1# set interface ge-0/2/0.0
user@PE1# set interface ge-0/2/4.0
```

4. Configure BGP.

```
[edit protocols bgp group pe-ce]
user@PE1# set type internal
user@PE1# set local-address 10.255.8.207
user@PE1# set family inet-vpn unicast
user@PE1# set neighbor 10.255.8.86
user@PE1# set export send-routes
```

5. Configure a policy that advertises direct and local interface routes.

```
[edit policy-options policy-statement send-routes term 1]
user@PE1# set from protocol direct
user@PE1# set from protocol local
user@PE1# set then accept
```

6. Configure an interior gateway protocol, such as IS-IS or OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface ge-0/2/0.0
user@PE1# set interface ge-0/2/4.0
user@PE1# set interface lo0.0 passive
```

7. Configure a signaling protocol, such as RSVP or LDP.

```
[edit protocols ldp]
user@PE1# set interface ge-0/2/0.0
user@PE1# set interface ge-0/2/4.0
```

8. (Optional) Configure nonstop active routing.

```
[edit routing-options]
user@PE1# set nonstop-routing
```

9. Configure the autonomous system (AS).

```
[edit routing-options]
user@PE1# set routing-options autonomous-system 100
```

10. Configure the Layer 3 VPN routing instance.

```
[edit routing-instances cust1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-4/1/0.0
user@PE1# set route-distinguisher 100:100
user@PE1# set vrf-target target:100:100
user@PE1# set vrf-table-label
```

11. Configure HFRR link protection.

```
[edit routing-instances cust1 routing-options]
user@PE1# set interface ge-4/1/0.0 link-protection (Host Fast Reroute)
```

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

```
[edit]
```

```
user@PE1# commit
```

### Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, **show policy-options**, **show routing-options**, and **show routing-instances** commands.

```
user@PE1# show interfaces
ge-4/1/0 {
  unit 0 {
    description toPE2;
    family inet {
      address 192.0.2.2/24 {
        vrrp-group 1 {
          virtual-address 192.0.2.5;
          priority 240;
          fast-interval 100;
          preempt;
          accept-data;
        }
      }
    }
  }
}
ge-0/2/0 {
  unit 0 {
    description toP1;
    family inet {
      address 10.10.10.1/30;
    }
    family mpls;
  }
}
ge-0/2/4 {
  unit 0 {
    description toP5;
    family inet {
      address 10.10.15.2/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.8.207/32;
    }
  }
}

user@PE1# show protocols
mpls {
  interface ge-0/2/0.0;
  interface ge-0/2/4.0;
}
```

```
bgp {
  group pe-ce {
    export-send-routes;
    type internal;
    local-address 10.255.8.207;
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.8.86;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/2/0.0;
    interface ge-0/2/4.0;
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface ge-0/2/0.0;
  interface ge-0/2/4.0;
}

user@PE1# show policy-options
policy-statement send-routes {
  term 1 {
    from protocol [ direct local ];
    then accept;
  }
}

user@PE1# show routing-options
nonstop-routing;
autonomous-system 100;

user@PE1# show routing-instances
cust1 {
  instance-type vrf;
  interface ge-4/1/0.0;
  route-distinguisher 100:100;
  vrf-target target:100:100;
  vrf-table-label;
  routing-options {
    interface {
      ge-4/1/0.0 {
        link-protection;
      }
    }
  }
}
```



---

## Verification

---

Confirm that the configuration is working properly.

- [Verifying HFRR on page 925](#)
- [Verifying ARP Routes on page 925](#)
- [Verifying Fast Reroute Routes on page 926](#)
- [Verifying Forwarding on page 927](#)

### *Verifying HFRR*

**Purpose** Make sure that HFRR is enabled.

**Action** user@PE1> [show hfr profiles](#)  
HFRR pointer: 0x9250000  
HFRR Current State: HFRR\_ACTIVE  
HFRR Protected IFL Name: ge-4/1/0.0  
HFRR Protected IFL Handle: 0x921086c  
HFRR Routing Instance Name: cust1  
HFRR Routing Instance Handle: 0x9129740  
HFRR Sync BG Scheduled: NO  
HFRR RTS Filter On: YES  
HFRR Delete BG Scheduled: NO  
HFRR Num ARP Routes learnt: 100  
HFRR Num FRR Routes Created: 100

**Meaning** The output shows that the HFRR is enabled on interface ge-4/1/0.0.

### *Verifying ARP Routes*

**Purpose** Make sure that the expected ARP routes are learned.

**Action** user@PE1> 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

192.0.2.3/24      @[ARP/4294967293] 00:04:35, from 192.0.2.1
                  Unusable
192.0.2.4/24      @[ARP/4294967293] 00:04:35, from 192.0.2.1
                  Unusable
192.0.2.5/24      @[ARP/4294967293] 00:04:32, from 192.0.2.1
                  Unusable
192.0.2.6/24      @[ARP/4294967293] 00:04:34, from 192.0.2.1
                  Unusable
192.0.2.7/24      @[ARP/4294967293] 00:04:35, from 192.0.2.1
                  Unusable
192.0.2.8/24      @[ARP/4294967293] 00:04:35, from 192.0.2.1
                  Unusable
192.0.2.9/24      @[ARP/4294967293] 00:04:35, from 192.0.2.1
                  Unusable
192.0.2.10/24     @[ARP/4294967293] 00:04:35, from 192.0.2.1
                  Unusable
192.0.2.11/24     @[ARP/4294967293] 00:04:33, from 192.0.2.1
                  Unusable
192.0.2.12/24     @[ARP/4294967293] 00:04:33, from 192.0.2.1
                  Unusable
192.0.2.13/24     @[ARP/4294967293] 00:04:33, from 192.0.2.1
                  Unusable
...
```

### *Verifying Fast Reroute Routes*

**Purpose** Make sure that the expected fast reroute (FRR) routes are learned.

**Action** user@PE1> 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)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.3/24      #[FRR/200] 00:05:38, from 192.0.2.1
                  > to 192.0.2.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.4/24      #[FRR/200] 00:05:38, from 192.0.2.1
                  > to 192.0.2.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.5/24      #[FRR/200] 00:05:35, from 192.0.2.1
                  > to 192.0.2.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.6/24      #[FRR/200] 00:05:37, from 192.0.2.1
                  > to 192.0.2.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.7/24      #[FRR/200] 00:05:38, from 192.0.2.1
                  > to 192.0.2.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.8/24      #[FRR/200] 00:05:38, from 192.0.2.1
                  > to 192.0.2.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.9/24      #[FRR/200] 00:05:38, from 192.0.2.1
                  > to 192.0.2.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
192.0.2.10/24     #[FRR/200] 00:05:38, from 192.0.2.1
...

```

### *Verifying Forwarding*

**Purpose** Make sure that the expected routes appear in the forwarding table.

```

Action user@PE1> show route forwarding-table destination 192.0.2.3
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
rjct             36      1

Routing table: default-switch.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
rjct             554     1

Routing table: __master.anon__.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
rjct             532     1

Routing table: cust1.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
192.0.2.3/24      user  0
                                0:0:14:14:1:3    ucst  767      3 ge-4/1/0.0
                                10.10.15.1       Push 16, Push 299792(top) 1262
2 ge-0/2/4.0
192.0.2.3/24      dest  0 0:0:14:14:1:3    ucst  767      3 ge-4/1/0.0

...

```

**See Also** • [Understanding Host Fast Reroute on page 911](#)

**Related Documentation** • [Fast Reroute Overview](#)

## Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs

For Layer 3 VPNs configured on Juniper Networks routers, Junos OS normally allocates one inner VPN label for each customer edge (CE)-facing virtual routing and forwarding (VRF) interface of a provider edge (PE) router. However, other vendors allocate one VPN label for each route learned over the CE-facing interfaces of a PE router. This practice increases the number of VPN labels exponentially, which leads to slow system processing and slow convergence time.

Next-hop chaining (also known as “chained composite next hop”) is a composition function that concatenates the partial rewrite strings associated with individual next hops to form a larger rewrite string that is added to a packet. By using this function, the number of routes with unique inner VPN labels that can be processed by a Juniper Networks router is increased substantially. Common route update elements associated with Layer 3 VPNs are combined, reducing the number of route updates and individual states the Juniper Networks router must maintain, and leading to enhanced scaling and convergence performance.



**NOTE:** ACX Series routers supports chained-composite-next-hop ingress CLI statement at the [edit routing-options forwarding-table] hierarchy level only for Layer 3 VPNs. The chained-composite-next-hop ingress CLI statement for Layer 2 services is not supported.

You can configure the router based on the number of VPN labels you want to manage and on whether or not you want to create chained composite next hops for IPv6 labeled routes:

- [Accepting Up to One Million Layer 3 VPN Route Updates on page 929](#)
- [Accepting More Than One Million Layer 3 VPN Route Updates on page 930](#)
- [Enabling Chained Composite Next Hops for IPv6 Labeled Unicast Routes on page 932](#)

## Accepting Up to One Million Layer 3 VPN Route Updates

For Juniper Networks routers participating in a mixed vendor network with up to one million Layer 3 VPN labels, include the **l3vpn** statement at the [edit routing-options forwarding-table chained-composite-next-hop ingress] hierarchy level. The **l3vpn** statement is disabled by default.



**NOTE:** ACX Series routers do not support chained-composite-next-hop ingress CLI statement at the [edit routing-options forwarding-table] hierarchy level.



**BEST PRACTICE:** We recommend that you configure the **l3vpn** statement whenever you have deployed Juniper Networks routers in mixed vendor networks of up to one million routes to support Layer 3 VPNs.

Because using this statement can also enhance the Layer 3 VPN performance of Juniper Networks routers in networks where only Juniper Networks routers are deployed, we recommend configuring the statements in these networks as well.

You can configure the **l3vpn** statement on the following routers:

- ACX Series routers
- MX Series routers
- M120 routers
- M320 routers with one or more Enhanced III FPCs
- T Series routers (for Junos OS Release 10.4 and later)

To accept up to one million Layer 3 VPN route updates with unique inner VPN labels, configure the **l3vpn** statement. This statement is supported on indirectly connected PE routers only. Configuring this statement on a router that is directly connected to a PE router provides no benefit. You can configure the **l3vpn** statement on a router with a mix of links to both directly connected and indirectly connected PE routers.



**NOTE:** You cannot configure the **l3vpn** statement and sub-statements at same time that you have configured the **vpn-unequal-cost** statement.

To configure the router to accept up to one million Layer 3 VPN route updates with unique inner VPN labels:

1. Include the **l3vpn** statement.

```
[edit routing-options forwarding-table chained-composite-next-hop ingress]
user@host>set l3vpn
```

2. To enhance memory allocation to support a larger number of Layer 3 VPN labels, include the **vpn-label** statement.

```
[edit chassis memory-enhanced]
user@host>set vpn-label
```



**NOTE:** The **vpn-label** statement does not provide any functional changes when used on the MX Series routers. You can omit the configuration of this statement on MX Series routers.

For more information about configuring more memory for Layer 3 VPN labels, see the *Junos OS Administration Library*.

After you have configured the **l3vpn** statement, you can determine whether or not a Layer 3 VPN route is a part of a composite next hop by examining the display output of the following commands:

- **show route route-value extensive**
- **show route forwarding-table destination destination-value extensive**

## Accepting More Than One Million Layer 3 VPN Route Updates

For Juniper Networks routers participating in a mixed vendor network with more than one million Layer 3 VPN labels, include the **extended-space** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]** hierarchy level. The **extended-space** statement is disabled by default.



**NOTE:** The **chained-composite-next-hop ingress** and **extended-space** statements are not supported on ACX Series routers.



**BEST PRACTICE:** We recommend that you configure the **extended-space** statement in mixed vendor networks containing more than one million routes to support Layer 3 VPNs. However, b

Because using this statements can also enhance the Layer 3 VPN performance of Juniper Networks routers in networks where only Juniper Networks routers are deployed, we recommend configuring the statement in these networks as well.

Using the **extended-space** statement can double the number of routes with unique inner VPN labels that can be processed by a Juniper Networks router. However, when configuring such very large-scale Layer 3 VPN scenarios, keep the following guidelines in mind:

- The **extended-space** statement is supported only on MX Series routers containing only MPCs.
- The chassis must be configured to use the **enhanced-ip** option in network services mode.

For more information about configuring chassis network services, see the *Junos OS Administration Library*.

- Ensure that you configure per-packet load balancing for associated policies.

For more information about configuring policies, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.



**BEST PRACTICE:** We strongly recommend using 64-bit routing engines running 64-bit Junos OS to support Layer 3 VPN prefixes with unique inner VPN labels at higher scale.

To configure the router to accept more than one million Layer 3 VPN route updates with unique inner VPN labels:

1. Include the **l3vpn** statement.

```
[edit routing-options forwarding-table chained-composite-next-hop ingress]
user@host> set l3vpn
```

2. Include the **extended-space** statement.

```
[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]
user@host> set extended-space
```

3. Configure chassis network services for enhanced mode.

```
[edit chassis]
user@host> set network-services enhanced-ip
```



**NOTE:** A router reboot might be required. See *Network Services Mode Overview* in the *Junos OS Administration Library* for details.

After you have completed the configuration, you can determine whether or not a Layer 3 VPN route is a part of a composite next hop by examining the display output of the following commands:

- **show route *route-value* extensive**
- **show route forwarding-table destination *destination-value* extensive**

### Enabling Chained Composite Next Hops for IPv6 Labeled Unicast Routes

You can enable chained composite next hops for IPv6 labeled unicast routes by configuring the **labeled-bgp** and **inet6** statements:

- To enable chained composite next hops for inet6 labeled unicast routes, include the **inet6** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp]** hierarchy level. This statement is disabled by default.

#### Related Documentation

- *Chained Composite Next Hops for Transit Devices for VPNs*
- *Configuring the Junos OS to Allocate More Memory for Routing Tables, Firewall Filters, and Layer 3 VPN Labels*
- *Network Services Mode Overview*
- *Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers*



## PART 9

# Connecting Layer 3 VPNs to Other VPNs

- [Connecting Layer 3 VPNs to Layer 2 Circuits on page 935](#)
- [Connecting Layer 3 VPNs to Layer 2 VPNs on page 957](#)



# Connecting Layer 3 VPNs to Layer 2 Circuits

- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 935](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 936](#)

## Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN

---

MPLS-based Layer 2 services are growing in demand among enterprise and service providers. This creates new challenges related to interoperability between Layer 2 and Layer 3 services for service providers who want to provide end-to-end value-added services. There are various reasons to stitch different Layer 2 services to one another and to Layer 3 services. For example, to expand the service offerings and to expand geographically. The Junos OS has various features to address the needs of the service provider.

You can enable pseudowire services and configure a pseudowire service interface as an access point for interconnecting layer 2 circuits to layer 3 VPNs. For more information, see *Pseudowire Subscriber Logical Interfaces Overview*.

Interconnecting a Layer 2 Circuit with a Layer 3 VPN provides the following benefits:

- Interconnecting a Layer 2 Circuit with a Layer 3 VPN enables the sharing of a service provider's core network infrastructure between IP and Layer 2 circuit services, reducing the cost of providing those services. A Layer 2 MPLS circuit allows service providers to create a Layer 2 circuit service over an existing IP and MPLS backbone.
- Service providers do not have to invest in separate Layer 2 equipment to provide Layer 2 circuit service. A service provider can configure a provider edge router to run any Layer 3 protocol in addition to the Layer 2 protocols. Customers who prefer to maintain control over most of the administration of their own networks want Layer 2 circuit connections with their service provider instead of a Layer 3 VPN connection.

### Related Documentation

- [Understanding Layer 3 VPNs on page 123](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 936](#)
- [Pseudowire Subscriber Logical Interfaces Overview](#)
- [Configuring a Pseudowire Subscriber Logical Interface](#)

## Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN

This example provides a step-by-step procedure and commands for configuring and verifying a Layer 2 circuit to Layer 3 VPN interconnection. It contains the following sections:

- [Requirements on page 936](#)
- [Overview and Topology on page 936](#)
- [Configuration on page 937](#)
- [Verifying the Layer 2 Circuit to Layer 3 VPN Interconnection on page 948](#)

### Requirements

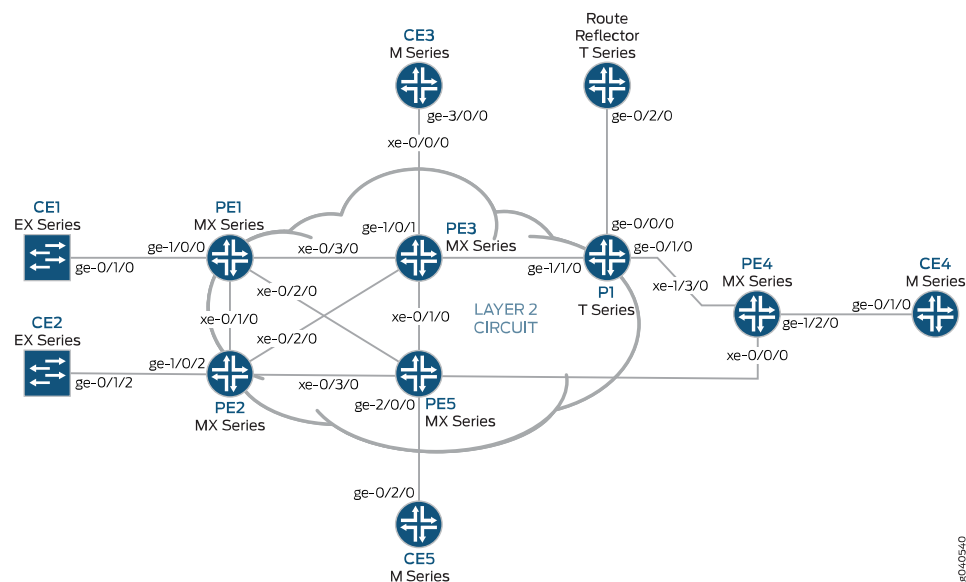
This example uses the following hardware and software components:

- Junos OS Release 9.3 or later
- 3 MX Series 3D Universal Edge Routers
- 1 M Series Multiservice Edge Router
- 1 T Series Core Router
- 1 EX Series Ethernet Switch

### Overview and Topology

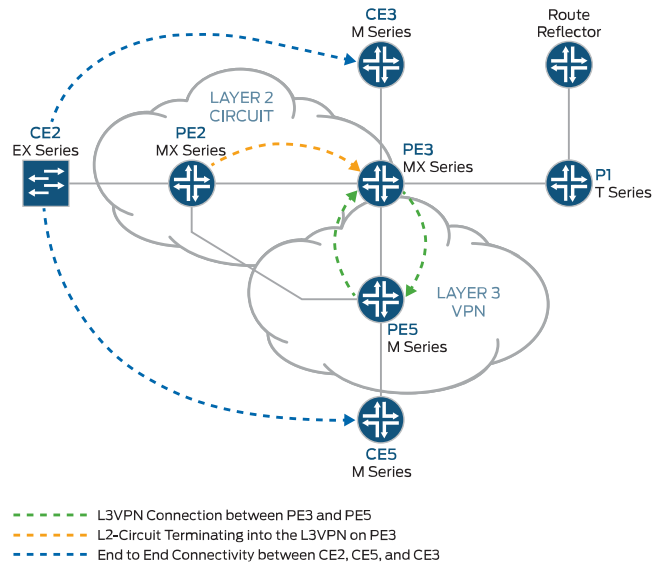
The physical topology of a Layer 2 circuit to Layer 3 VPN interconnection is shown in [Figure 86 on page 936](#).

**Figure 86: Physical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection**



The logical topology of a Layer 2 circuit to Layer 3 VPN interconnection is shown in [Figure 87 on page 937](#).

**Figure 87: Logical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection**



## Configuration



**NOTE:** In any configuration session, it is good practice to verify periodically that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE2** identifies the customer edge 2 (CE2) router
- **PE1** identifies the provider edge 1 (PE1) router
- **CE3** identifies the customer edge 3 (CE3) router
- **PE3** identifies the provider edge 3 (PE3) router
- **CE5** identifies the customer edge 5 (CE5) router
- **PE5** identifies the provider edge 5 (PE5) router

This example contains the following procedures:

- [Configuring PE Router Customer-facing and Loopback Interfaces on page 938](#)
- [Configuring Core-facing Interfaces on page 939](#)
- [Configuring Protocols on page 941](#)
- [Configuring Routing Instances and Layer 2 Circuits on page 944](#)

- [Configuring the Route Reflector on page 945](#)
- [Interconnecting the Layer 2 Circuit with the Layer 3 VPN on page 947](#)

### Configuring PE Router Customer-facing and Loopback Interfaces

#### Step-by-Step Procedure

To begin building the interconnection, configure the interfaces on the PE routers. If your network contains provider (P) routers, configure the interfaces on the P routers also. This example shows the configuration for Router PE2, Router PE3, and Router PE5.

1. On Router PE2, configure the **ge-1/0/2** interface encapsulation. To configure the interface encapsulation, include the **encapsulation** statement and specify the **ethernet-ccc** option (**vlan-ccc** encapsulation is also supported). Configure the **ge-1/0/2.0** logical interface family for circuit cross-connect functionality. To configure the logical interface family, include the **family** statement and specify the **ccc** option. The encapsulation should be configured the same way for all routers in the Layer 2 circuit domain.

```
[edit interfaces]
ge-1/0/2 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
```

2. On Router PE2, configure the **lo0.0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **192.0.2.2/24** as the loopback IPv4 address.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
}
```

3. On Router PE3, configure the **ge-1/0/1** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **198.51.100.1/24** as the interface address for this device.

```
[edit interfaces]
ge-1/0/1 {
  unit 0 {
    family inet {
      address 198.51.100.1/24;
    }
  }
}
```

4. On Router PE3, configure the **lo0.0** loopback interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **192.0.2.3/24** as the loopback IPv4 address for this router.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.3/24;
    }
  }
}
```

5. On Router PE5, configure the **ge-2/0/0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **198.51.100.8/24** as the interface address.

```
[edit interfaces]
ge-2/0/0 {
  unit 0 {
    family inet {
      address 198.51.100.8/24;
    }
  }
}
```

6. On Router PE5, configure the **lo0.0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **192.0.2.5/24** as the loopback IPv4 address for this router.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 192.0.2.5/24;
    }
  }
}
```

### Configuring Core-facing Interfaces

#### Step-by-Step Procedure

This procedure describes how to configure the core-facing interfaces on the PE routers. This example does not include all the core-facing interfaces shown in the physical topology illustration. Enable the **mpls** and **inet** address families on the core-facing interfaces.

1. On Router PE2, configure the **xe-0/2/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.5.1/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```
[edit interfaces]
```

```
xe-0/2/0 {  
  unit 0 {  
    family inet {  
      address 10.10.5.1/30;  
    }  
    family mpls;  
  }  
}
```

2. On Router PE3, configure the core-facing interfaces. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify the IPv4 addresses shown in the example as the interface addresses. Include the **family** statement and specify the **mpls** address family. In the example, the **xe-2/1/0** interface is connected to Router PE5, and the **xe-2/2/0** interface is connected to Router PE2.

```
[edit interfaces]  
xe-2/0/0 {  
  unit 0 {  
    family inet {  
      address 10.10.20.2/30;  
    }  
    family mpls;  
  }  
}  
xe-2/1/0 {  
  unit 0 {  
    family inet {  
      address 10.10.6.1/30;  
    }  
    family mpls;  
  }  
}  
xe-2/2/0 {  
  unit 0 {  
    family inet {  
      address 10.10.5.2/30;  
    }  
    family mpls;  
  }  
}  
xe-2/3/0 {  
  unit 0 {  
    family inet {  
      address 10.10.1.2/30;  
    }  
    family mpls;  
  }  
}
```

3. On Router PE5, configure the **xe-0/1/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.6.2/30** as the interface address. Include the **family** statement and specify the **mpls** address family.



```
[edit interfaces]
xe-0/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.2/30;
    }
    family mpls;
  }
}
```

### Configuring Protocols

#### Step-by-Step Procedure

This procedure describes how to configure the protocols used in this example. If your network contains P routers, configure the interfaces on the P routers also.

1. On Router PE3, enable OSPF as the IGP. Enable the MPLS, LDP, and BGP protocols on all interfaces except **fxp0.0**. LDP is used as the signaling protocol for the Layer 2 circuit to Router PE2. The following configuration snippet shows the protocol configuration for Router PE3:

```
[edit]
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-RR {
      to 192.0.2.7;
    }
    label-switched-path to-PE2 {
      to 192.0.2.2;
    }
    label-switched-path to-PE5 {
      to 192.0.2.5;
    }
    label-switched-path to-PE4 {
      to 192.0.2.4;
    }
    label-switched-path to-PE1 {
      to 192.0.2.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 192.0.2.3;
      family inet-vpn {
```

```
        unicast;
      }
      family l2vpn {
        signaling;
      }
      neighbor 192.0.2.7;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
ldp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
}
```

2. On Router PE2, configure the MPLS, OSPF, and LDP protocols.

```
[edit ]
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
ldp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
}
```

3. On Router PE5, enable OSPF as the IGP. Enable the MPLS, RSVP, and BGP protocols on all interfaces except **fxp0.0**. Enable core-facing interfaces with the **mpls** and **inet** address families.

```
[edit]
protocols {
  rsvp {
    interface all {
      link-protection;
    }
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-RR {
      to 192.0.2.7;
    }
    label-switched-path to-PE2 {
      to 192.0.2.2;
    }
    label-switched-path to-PE3 {
      to 192.0.2.3;
    }
    label-switched-path to-PE4 {
      to 192.0.2.4;
    }
    label-switched-path to-PE1 {
      to 192.0.2.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group to-rr {
      type internal;
      local-address 192.0.2.5;
      family inet-vpn {
        unicast;
      }
      family l2vpn {
        signaling;
      }
      neighbor 192.0.2.7;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

```
}
}
```

## Configuring Routing Instances and Layer 2 Circuits

### Step-by-Step Procedure

This procedure describes how to configure the Layer 2 circuit and the Layer 3 VPN.

1. On Router PE2, configure the Layer 2 circuit. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE3 as the neighbor. Include the interface statement and specify **ge-1/0/2.0** as the logical interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **no-control-word** statement for equipment that does not support the control word.

```
[edit ]
protocols {
  l2circuit {
    neighbor 192.0.2.3 {
      interface ge-1/0/2.0 {
        virtual-circuit-id 100;
        no-control-word;
      }
    }
  }
}
```

2. On Router PE3, configure the Layer 2 circuit to Router PE2. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE2 as the neighbor. Include the interface statement and specify **lt-1/1/10.0** as the logical tunnel interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **no-control-word** statement.

```
[edit ]
protocols {
  l2circuit {
    neighbor 192.0.2.2 {
      interface lt-1/1/10.0 {
        virtual-circuit-id 100;
        no-control-word;
      }
    }
  }
}
```

3. On Router PE3, configure the Layer 3 VPN (**L3VPN**) routing instance to Router PE5 at the **[edit routing-instances]** hierarchy level. Also configure the BGP peer group at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```
[edit ]
routing-instances {
  L3VPN {
```

```

instance-type vrf;
interface ge-1/0/1.0;
interface lt-1/1/10.1;
route-distinguisher 65000:33;
vrf-target target:65000:2;
vrf-table-label;
protocols {
  bgp {
    export direct;
    group ce3 {
      neighbor 198.51.100.6 {
        peer-as 100;
      }
    }
  }
}
}
}

```

4. On Router PE5, configure the Layer 3 VPN routing instance (**L3VPN**) at the **[edit routing-instances]** hierarchy level. Also configure the BGP peer group at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```

[edit ]
routing-instances {
  L3VPN {
    instance-type vrf;
    interface ge-2/0/0.0;
    route-distinguisher 65000:5;
    vrf-target target:65000:2;
    vrf-table-label;
    protocols {
      bgp {
        group ce5 {
          neighbor 198.51.100.10 {
            peer-as 200;
          }
        }
      }
    }
  }
}
}

```

### Configuring the Route Reflector

#### Step-by-Step Procedure

Although a route reflector is not required to interconnect a Layer 2 circuit with a Layer 3 VPN, this examples uses a route reflector. This procedure shows the relevant portion of the route reflector configuration.

1. Configure the route reflector with RSVP, MPLS, BGP and OSPF. The route reflector is a BGP peer with the PE routers. Notice that the BGP peer group configuration includes the **family** statement and specifies the **inet-vpn** option The **inet-vpn** option enables BGP to advertise network layer reachability information (NLRI) for the Layer

3 VPN routes. The configuration also includes the **family** statement and specifies the **l2vpn** option. The **l2vpn** option enables BGP to advertise NLRI for the Layer 2 circuit. Layer 2 circuits use the same internal BGP infrastructure as Layer 2 VPNs.

```
[edit ]
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-pe3 {
      to 192.0.2.3;
    }
    label-switched-path to-pe5 {
      to 192.0.2.5;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 192.0.2.7;
      family inet {
        unicast;
      }
      family inet-vpn {
        unicast;
      }
      family l2vpn {
        signaling;
      }
      cluster 192.0.2.7;
      neighbor 192.0.2.1;
      neighbor 192.0.2.2;
      neighbor 192.0.2.4;
      neighbor 192.0.2.5;
      neighbor 192.0.2.3;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

### Interconnecting the Layer 2 Circuit with the Layer 3 VPN

#### Step-by-Step Procedure

Before you can configure the logical tunnel interface in an MX Series router, you must create the tunnel services interface to be used for tunnel services.

1. Create the tunnel service interface on Router PE3. Include the **bandwidth** statement at the **[edit chassis fpc slot-number pic slot-number tunnel-services]** hierarchy level and specify the amount of bandwidth to reserve for tunnel services in gigabits per second.

```
[edit chassis]
fpc 1 {
  pic 1 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
```

2. On Router PE3, configure the **lt-1/1/10** logical tunnel interface unit 0.

Router PE3 is the router that is *stitching* the Layer 2 circuit to the Layer 3 VPN using the logical tunnel interface. The configuration of the peer unit interfaces is what makes the interconnection.

Include the **encapsulation** statement and specify the **ethernet-ccc** option. Include the **peer-unit** statement and specify the logical interface unit 1 as the peer tunnel interface. Include the **family** statement and specify the **ccc** option.

Configure the **lt-1/1/10** logical interface unit 1 with **ethernet** encapsulation. Include the **peer-unit** statement and specify the logical interface unit 0 as the peer tunnel interface. Include the **family** statement and specify the **inet** option. Also include the **address** statement and specify **198.51.100.11/24** as the IPv4 address of the interface.



**NOTE:** The peering logical interfaces must belong to the same logical tunnel interface derived from the Tunnel Services PIC.

```
[edit interfaces]
lt-1/1/10 {
  unit 0 {
    encapsulation ethernet-ccc;
    peer-unit 1;
    family ccc;
  }
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet {
      address 198.51.100.11/24;
    }
  }
}
```

```
}
```

- On each router, commit the configuration.

```
user@host> commit check
configuration check succeeds
user@host> commit
```

## Verifying the Layer 2 Circuit to Layer 3 VPN Interconnection

To verify that the interconnection is working properly, perform these tasks:

- [Verifying That the Layer 2 Circuit Connection to Router PE3 is Up on page 948](#)
- [Verifying LDP Neighbors and Targeted LDP LSPs on Router PE2 on page 949](#)
- [Verifying the Layer 2 Circuit Routes on Router PE2 on page 949](#)
- [Verifying That the Layer 2 Circuit Connection to Router PE2 is Up on page 950](#)
- [Verifying LDP Neighbors and Targeted LDP LSPs on Router PE3 on page 951](#)
- [Verifying a BGP Peer Session with the Route Reflector on Router PE3 on page 951](#)
- [Verifying the Layer 3 VPN Routes on Router PE3 on page 952](#)
- [Verifying the Layer 2 Circuit Routes on Router PE3 on page 952](#)
- [Verifying the MPLS Routes on Router PE3 on page 953](#)
- [Verifying Traffic Flow Between Router CE2 and Router CE3 on page 954](#)
- [Verifying Traffic Flow Between Router CE2 and Router CE5 on page 954](#)

### Verifying That the Layer 2 Circuit Connection to Router PE3 is Up

**Purpose** To verify that the Layer 2 circuit connection from Router PE2 to Router PE3 is **Up**. To also document the incoming and outgoing LDP labels and the circuit ID used by this Layer 2 circuit connection.

**Action** Verify that the Layer 2 circuit connection is up, using the **show l2circuit connections** command.

```
user@PE2> show l2circuit connections
```

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	SP -- Static Pseudowire
LD -- local site signaled down	RS -- remote site standby
RD -- remote site signaled down	XX -- unknown

Legend for interface status



```

Up -- operational
Dn -- down
Neighbor: 192.0.2.3
Interface          Type St    Time last up      # Up trans
ge-1/0/2.0(vc 100) rmt  Up    Jan 7 02:14:13 2010      1
Remote PE: 192.0.2.3, Negotiated control-word: No
Incoming label: 301488, Outgoing label: 315264
Negotiated PW status TLV: No
Local interface: ge-1/0/2.0, Status: Up, Encapsulation: ETHERNET

```

**Meaning** The output shows that the Layer 2 circuit connection from Router PE2 to Router PE3 is **Up** and the connection is using the **ge-1/0/2.0** interface. Note that the outgoing label is **315264** and the incoming label is **301488**, the virtual circuit (VC) identifier is **100** and the encapsulation is **ETHERNET**.

### Verifying LDP Neighbors and Targeted LDP LSPs on Router PE2

**Purpose** To verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors.

**Action** Verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors, using the **show ldp neighbor** command.

```

user@PE2> show ldp neighbor
Address          Interface      Label space ID      Hold time
192.0.2.3        lo0.0          192.0.2.3:0         38

```

**Meaning** The output shows that Router PE2 has an LDP neighbor with the IPv4 address of **192.0.2.3**. Address 192.0.2.3 is the lo0.0 interface address of Router PE3. Notice that Router PE2 uses the local **lo0.0** interface for the LSP.

Verifying that the routers are LDP neighbors also verifies that the targeted LSP is established.

### Verifying the Layer 2 Circuit Routes on Router PE2

**Purpose** To verify that Router PE2 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE3.

**Action** Verify that Router PE2 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE3, using the **show route table mpls.0** command.

```

user@PE2> show route table mpls.0
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 1w3d 05:24:11, metric 1
                  Receive

```

```

1          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
2          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
300560     *[LDP/9] 16:12:23, metric 1
           > to 10.10.2.1 via xe-0/1/0.0, Pop
300560(S=0) *[LDP/9] 16:12:23, metric 1
           > to 10.10.2.1 via xe-0/1/0.0, Pop
301008     *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Swap 299856
301488     *[L2CKT/7] 11:07:28
           > via ge-1/0/2.0, Pop
301536     *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Pop
301536(S=0) *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Pop
301712     *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Swap 315184
301728     *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Pop
301728(S=0) *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Pop
ge-1/0/2.0 *[L2CKT/7] 11:07:28, metric2 1
           > to 10.10.5.2 via xe-0/2/0.0, Push 315264

```

**Meaning** The output shows that Router PE2 pushes the **315264** outgoing label on the **L2CKT** route going out interface **ge-1/0/2.0**. The output also shows that Router PE2 pops the **301488** incoming label on the **L2CKT** coming from interface **ge-1/0/2.0**

### Verifying That the Layer 2 Circuit Connection to Router PE2 is Up

**Purpose** To verify that the Layer 2 circuit connection from Router PE3 to Router PE2 is **Up**, To also document the incoming and outgoing LDP labels and the circuit ID used by this Layer 2 circuit connection.

**Action** Verify that the Layer 2 circuit connection is up, using the **show l2circuit connections** command.

user@PE3> show l2circuit connections

Layer-2 Circuit Connections:

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	XX -- unknown

Legend for interface status

Up -- operational  
Dn -- down

```
Neighbor: 192.0.2.2
Interface                Type St   Time last up      # Up trans
lt-1/1/10.0(vc 100)      rmt  Up   Jan  7 02:15:03 2010      1
Remote PE: 192.0.2.2, Negotiated control-word: No
Incoming label: 315264, Outgoing label: 301488
Local interface: lt-1/1/10.0, Status: Up, Encapsulation: ETHERNET
```

**Meaning** The output shows that the Layer 2 circuit connection from Router PE3 to Router PE2 is **Up** and the connection is using the logical tunnel (**lt**) interface. Note that the incoming label is **315264** and the outgoing label is **301488**, the virtual circuit (VC) identifier is **100**, and that the encapsulation is **ETHERNET**.

### Verifying LDP Neighbors and Targeted LDP LSPs on Router PE3

**Purpose** To verify that Router PE3 has a targeted LDP LSP to Router PE2 and that Router PE3 and Router PE2 are LDP neighbors.

**Action** Verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors, using the **show ldp neighbor** command.

```
user@PE2> show ldp neighbor
Address      Interface      Label space ID      Hold time
192.0.2.2    lo0.0          192.0.2.2:0         43
192.0.2.4    lo0.0          192.0.2.4:0         33
```

**Meaning** The output shows that Router PE3 has an LDP neighbor with the IPv4 address of **192.0.2.2**. Address 192.0.2.2 is the lo0.0 interface address of Router PE2. The output also shows that the interface used on Router PE3 for the LSP is **lo0.0**. Verifying that the routers are LDP neighbors also verifies that the targeted LSP is established.

### Verifying a BGP Peer Session with the Route Reflector on Router PE3

**Purpose** To verify that Router PE3 has a peer session established with the route reflector.

**Action** Verify that Router PE3 has a peer session established with the route reflector, using the **show bgp summary** command.

```
user@PE2> show bgp summary

Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.13vpn.0      1          1          0          0          0          0
Peer          AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.0.2.7      65000     1597     1612      0        1   12:03:21 Establ
  bgp.12vpn.0: 0/0/0/0
  bgp.13vpn.0: 1/1/1/0
  L3VPN.inet.0: 1/1/1/0
```

**Meaning** The output shows that Router PE3 has a peer session with the router with the IPv4 address of **192.0.2.7**. Address 192.0.2.7 is the lo0.0 interface address of the route reflector. The output also shows that the peer session state is **Establ**, meaning that the session is established.

### Verifying the Layer 3 VPN Routes on Router PE3

**Purpose** To verify that Router PE3 has Layer 3 VPN routes to Router CE2, Router CE3, and Router CE5.

**Action** Verify that Router PE3 has routes to Router CE2, Router CE3, and Router CE5 in the Layer 3 VPN route table, using the **show route table L3VPN.inet.0** command. In this example, **L3VPN** is the name configured for the routing instance.

```
user@PE3> show route table L3VPN.inet.0
L3VPN.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.10/24      *[Direct/0] 11:13:59
> via lt-1/1/10.1
198.51.100.11/24      *[Local/0] 11:13:59
Local via lt-1/1/10.1
198.51.100.12/24      *[BGP/170] 11:00:41, localpref 100, from 192.0.2.7
AS path: I
> to 10.10.6.2 via xe-2/1/0.0, Push 16
198.51.100.13/24      *[Direct/0] 11:54:41
> via ge-1/0/1.0
198.51.100.1/24       *[Local/0] 11:54:41
Local via ge-1/0/1.0
```

**Meaning** The output shows that Router PE3 has a route to the IPv4 subnetwork address of **198.51.100.10**. Address 198.51.100.15 is the interface address of Router CE2. The output shows that Router PE3 has a route to the IPv4 subnetwork address of **198.51.100.12**. Address 198.51.100.10 is the interface address of Router CE5. The output shows that Router PE3 has a route to the IPv4 subnetwork address of **198.51.100.13**. Address 198.51.100.6 is the interface address of Router CE3.

### Verifying the Layer 2 Circuit Routes on Router PE3

**Purpose** To verify that Router PE3 has a route to Router PE2 in the Layer 2 circuit route table.

**Action** Verify that Router PE3 has a route to Router PE2 in the Layer 2 circuit route table, using the **show route table l2circuit.0** command.

```
user@PE3> show route table l2circuit.0
192.0.2.2:NoCtrlWord:5:100:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop type: Indirect
    Next-hop reference count: 1
    Next hop type: Router
```

```

Next hop: 10.10.5.1 via xe-2/2/0.0, selected
Protocol next hop: 192.0.2.2
Indirect next hop: 8cae0a0 -
State: <Active Int>
Local AS: 65000
Age: 11:16:50 Metric2: 1
Task: 12 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 315264, MTU 1500

```

**Meaning** The output shows that Router PE3 has a route to the IPv4 address of **192.0.2.2**. Address 192.0.2.2 is the lo0.0 interface address of Router PE2. Note that the VC label is **315264**. This label is the same as the incoming MPLS label displayed using the **show l2circuit connections** command.

### Verifying the MPLS Routes on Router PE3

**Purpose** To verify that Router PE3 has a route to Router PE2 in the MPLS route table.

**Action** Verify Router PE3 has a route to Router PE2 in the MPLS route table, using the **show route table mpls.0** command.

```

user@PE3> show route table mpls.0
mpls.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
1          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
2          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
16         *[VPN/0] 12:22:45
            to table L3VPN.inet.0, Pop
315184     *[LDP/9] 12:45:14, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Pop
315184(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Pop
315200     *[LDP/9] 00:03:53, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Swap 625297
            to 10.10.6.2 via xe-2/1/0.0, Swap 299856
315216     *[LDP/9] 12:45:14, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, Pop
315216(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, Pop
315232     *[LDP/9] 12:45:06, metric 1
            > to 10.10.1.1 via xe-2/3/0.0, Pop
315232(S=0) *[LDP/9] 12:45:06, metric 1
            > to 10.10.1.1 via xe-2/3/0.0, Pop
315248     *[LDP/9] 12:45:14, metric 1
            > to 10.10.5.1 via xe-2/2/0.0, Pop
315248(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.5.1 via xe-2/2/0.0, Pop
315264     *[L2CKT/7] 11:11:20

```

```

> via lt-1/1/10.0, Pop
315312      *[RSVP/7] 11:26:01, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315312(S=0) *[RSVP/7] 11:26:01, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315328      *[RSVP/7] 11:26:01, metric 1
> to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
315360      *[RSVP/7] 11:26:01, metric 1
> to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
316208      *[RSVP/7] 00:03:32, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
316208(S=0) *[RSVP/7] 00:03:32, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
lt-1/1/10.0 *[L2CKT/7] 11:11:20, metric2 1
> to 10.10.5.1 via xe-2/2/0.0, Push 301488

```

**Meaning** The output shows that Router PE3 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE2. Notice that the **301488** label is the same as the outgoing label displayed on Router PE2 using the **show l2circuit connections** command.

### Verifying Traffic Flow Between Router CE2 and Router CE3

**Purpose** To verify that the CE routers can send and receive traffic across the interconnection.

**Action** Verify that Router CE2 can send traffic to and receive traffic from Router CE3 across the interconnection, using the **ping** command.

```

user@CE2>ping 198.51.100.6
PING 198.51.100.6 (198.51.100.6): 56 data bytes
64 bytes from 198.51.100.6: icmp_seq=0 ttl=63 time=0.708 ms
64 bytes from 198.51.100.6: icmp_seq=1 ttl=63 time=0.610 ms

```

**Meaning** The output shows that Router CE2 can send an ICMP request to and receive a response from Router CE3 across the interconnection.

### Verifying Traffic Flow Between Router CE2 and Router CE5

**Purpose** To verify that the CE routers can send and receive traffic across the interconnection.

**Action** Verify that Router CE2 can send traffic to and receive traffic from Router CE5 across the interconnection, using the **ping** command.

```

user@CE2>ping 198.51.100.10
PING 198.51.100.10 (198.51.100.10): 56 data bytes
64 bytes from 198.51.100.10: icmp_seq=0 ttl=62 time=0.995 ms
64 bytes from 198.51.100.10: icmp_seq=1 ttl=62 time=1.005 ms

```

**Meaning** The output shows that Router CE2 can send an ICMP request to and receive a response from Router CE5 across the interconnection.

- Related Documentation**
- [Understanding Layer 3 VPNs on page 123](#)
  - [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 935](#)





# Connecting Layer 3 VPNs to Layer 2 VPNs

- [Interconnecting Layer 2 VPNs with Layer 3 VPNs Overview on page 957](#)
- [Example: Interconnecting a Layer 2 VPN with a Layer 3 VPN on page 958](#)

## Interconnecting Layer 2 VPNs with Layer 3 VPNs Overview

---

As MPLS-based Layer 2 services grow in demand, new challenges arise for service providers to be able to interoperate with Layer 2 and Layer 3 services and give their customers value-added services. Junos OS has various features to address the needs of service providers. One of these features is the use of a logical tunnel interface. This Junos OS functionality makes use of a tunnel PIC to loop packets out and back from the Packet Forwarding Engine to link the Layer 2 network with the Layer 3 network. The solution is limited by the logical tunnel bandwidth constraints imposed by the tunnel PIC.

## Interconnecting Layer 2 VPNs with Layer 3 VPNs Applications

Interconnecting a Layer 2 VPN with a Layer 3 VPN provides the following benefits:

- **A single access line to provide multiple services**—Traditional VPNs over Layer 2 circuits require the provisioning and maintenance of separate networks for IP and for VPN services. In contrast, Layer 2 VPNs enable the sharing of a provider's core network infrastructure between IP and Layer 2 VPN services, thereby reducing the cost of providing those services.
- **Flexibility**—Many different types of networks can be accommodated by the service provider. If all sites in a VPN are owned by the same enterprise, this is an intranet. If various sites are owned by different enterprises, the VPN is an extranet. A site can be located in more than one VPN.
- **Wide range of possible policies**—You can give every site in a VPN a different route to every other site, or you can force traffic between certain pairs of sites routed via a third site and so pass certain traffic through a firewall.
- **Scalable network**—This design enhances the scalability because it eliminates the need for provider edge (PE) routers to maintain all of the service provider's VPN routes. Each PE router maintains a VRF table for each of its directly connected sites. Each customer connection (such as a Frame Relay PVC, an ATM PVC, or a VLAN) is mapped to a specific VRF table. Thus, it is a port on the PE router and not a site that is associated with a VRF table. Multiple ports on a PE router can be associated with a single VRF

table. It is the ability of PE routers to maintain multiple forwarding tables that supports the per-VPN segregation of routing information.

- Use of route reflectors—Provider edge routers can maintain IBGP sessions to route reflectors as an alternative to a full mesh of IBGP sessions. Deploying multiple route reflectors enhances the scalability of the RFC 2547bis model because it eliminates the need for any single network component to maintain all VPN routes.
- Multiple VPNs are kept separate and distinct from each other—The customer edge routers do not peer with each other. Two sites have IP connectivity over the common backbone only, and only if there is a VPN which contains both sites. This feature keeps the VPNs separate and distinct from each other, even if two VPNs have an overlapping address space.
- Simple for customers to use—Customers can obtain IP backbone services from a service provider, and they do not need to maintain their own backbones.

**Related Documentation**

- [Understanding Layer 3 VPNs on page 123](#)
- [Example: Interconnecting a Layer 2 VPN with a Layer 3 VPN on page 958](#)

---

## Example: Interconnecting a Layer 2 VPN with a Layer 3 VPN

This example provides a step-by-step procedure and commands for interconnecting and verifying a Layer 2 VPN with a Layer 3 VPN. It contains the following sections:

- [Requirements on page 958](#)
- [Overview and Topology on page 958](#)
- [Configuration on page 961](#)
- [Verification on page 977](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.3 or later
- Five MX Series routers
- Three M Series routers
- Two T Series routers

### Overview and Topology

A Layer 2 VPN is a type of virtual private network (VPN) that uses MPLS labels to transport data. The communication occurs between the provider edge (PE) routers.

Layer 2 VPNs use BGP as the signaling protocol and, consequently, have a simpler design and require less provisioning overhead than traditional VPNs over Layer 2 circuits. BGP signaling also enables autodiscovery of Layer 2 VPN peers. Layer 2 VPNs can have either a full-mesh or a hub-and-spoke topology. The tunneling mechanism in the core network

is, typically, MPLS. However, Layer 2 VPNs can also use other tunneling protocols, such as GRE.

Layer 3 VPNs are based on RFC 2547bis, *BGP/MPLS IP VPNs*. RFC 2547bis defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone. RFC 2547bis VPNs are also known as BGP/MPLS VPNs because BGP is used to distribute VPN routing information across the provider's backbone, and MPLS is used to forward VPN traffic across the backbone to remote VPN sites.

Customer networks, because they are private, can use either public addresses or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the same private addresses used by other network users. MPLS/BGP VPNs solve this problem by adding a *route distinguisher*. A route distinguisher is a VPN identifier prefix that is added to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and within the Internet.

In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only. To separate a VPN's routes from routes in the public Internet or those in other VPNs, the PE router creates a separate routing table for each VPN called a VPN routing and forwarding (VRF) table. The PE router creates one VRF table for each VPN that has a connection to a customer edge (CE) router. Any customer or site that belongs to the VPN can access only the routes in the VRF tables for that VPN. Every VRF table has one or more extended community attributes associated with it that identify the route as belonging to a specific collection of routers. One of these, the *route target* attribute, identifies a collection of sites (VRF tables) to which a PE router distributes routes. The PE router uses the route target to constrain the import of remote routes into its VRF tables.

When an ingress PE router receives routes advertised from a directly connected CE router, it checks the received route against the VRF export policy for that VPN.

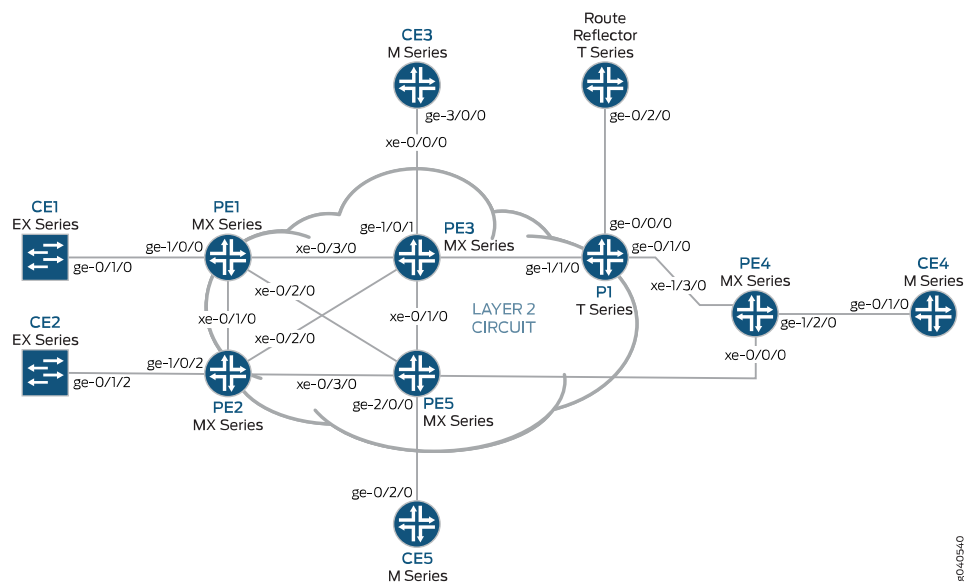
- If it matches, the route is converted to VPN-IPv4 format—that is, the route distinguisher is added to the route. The PE router then announces the route in VPN-IPv4 format to the remote PE routers. It also attaches a route target to each route learned from the directly connected sites. The route target attached to the route is based on the value of the VRF table's configured export target policy. The routes are then distributed using IBGP sessions, which are configured in the provider's core network.
- If the route from the CE router does not match, it is not exported to other PE routers, but it can still be used locally for routing, for example, if two CE routers in the same VPN are directly connected to the same PE router.

When an egress PE router receives a route, it checks it against the import policy on the IBGP session between the PE routers. If it is accepted, the router places the route into its `bgp.l3vpn.0` table. At the same time, the router checks the route against the VRF import policy for the VPN. If it matches, the route distinguisher is removed from the route and

the route is placed into the VRF table (the *routing-instance-name.inet.0* table) in IPv4 format.

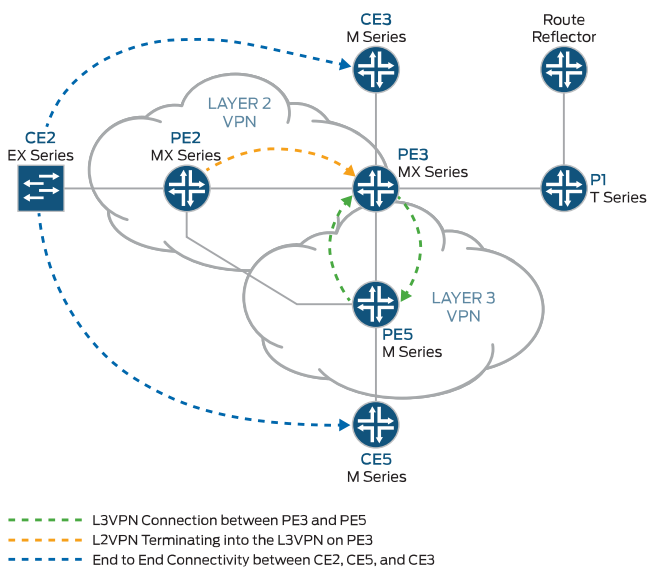
Figure 88 on page 960 shows the physical topology of a Layer 2 VPN-to-Layer 3 VPN interconnection.

Figure 88: Physical Topology of a Layer 2 VPN Terminating into a Layer 3 VPN



The logical topology of a Layer 2 VPN-to-Layer 3 VPN interconnection is shown in Figure 89 on page 960.

Figure 89: Logical Topology of a Layer 2 VPN Terminating into a Layer 3 VPN



The following definitions describe the meaning of the device abbreviations used in [Figure 88 on page 960](#) and [Figure 89 on page 960](#).

- Customer edge (CE) device—A device at the customer premises that provides access to the service provider's VPN over a data link to one or more provider edge (PE) routers.

Typically the CE device is an IP router that establishes an adjacency with its directly connected PE routers. After the adjacency is established, the CE router advertises the site's local VPN routes to the PE router and learns remote VPN routes from the PE router.

- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.

PE routers exchange routing information with CE routers. PE routers are aware of the VPNs that connect through them, and PE routers maintain VPN state. A PE router is only required to maintain VPN routes for those VPNs to which it is directly attached. After learning local VPN routes from CE routers, a PE router exchanges VPN routing information with other PE routers using IBGP. Finally, when using MPLS to forward VPN data traffic across the provider's backbone, the ingress PE router functions as the ingress label-switching router (LSR) and the egress PE router functions as the egress LSR.

- Provider (P) device—A device that operates inside the provider's core network and does not directly interface to any CE.

Although the P device is a key part of implementing VPNs for the service provider's customers and may provide routing for many provider-operated tunnels that belong to different VPNs, it is not itself VPN-aware and does not maintain VPN state. Its principal role is allowing the service provider to scale its VPN offerings, for example, by acting as an aggregation point for multiple PE routers.

P routers function as MPLS transit LSRs when forwarding VPN data traffic between PE routers. P routers are required only to maintain routes to the provider's PE routers; they are not required to maintain specific VPN routing information for each customer site.

## Configuration

To interconnect a Layer 2 VPN with a Layer 3 VPN, perform these tasks:

- [Configuring the Base Protocols and Interfaces on page 961](#)
- [Configuring the VPN Interfaces on page 965](#)

### Configuring the Base Protocols and Interfaces

#### Step-by-Step Procedure

1. On each PE and P router, configure OSPF with traffic engineering extensions on all interfaces. Disable OSPF on the fxp0.0 interface.

```
[edit protocols]
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
```

```
        interface fxp0.0 {  
            disable;  
        }  
    }  
}
```

2. On all the core routers, enable MPLS on all interfaces. Disable MPLS on the fxp0.0 interface.

```
[edit protocols]  
mpls {  
    interface all;  
    interface fxp0.0 {  
        disable;  
    }  
}
```

3. On all the core routers, create an internal BGP peer group and specify the route reflector address (192.0.2.7) as the neighbor. Also enable BGP to carry Layer 2 VPLS network layer reachability information (NLRI) messages for this peer group by including the **signaling** statement at the **[edit protocols bgp group group-name family l2vpn]** hierarchy level.

```
[edit protocols]  
bgp {  
    group RR {  
        type internal;  
        local-address 192.0.2.2;  
        family l2vpn {  
            signaling;  
        }  
        neighbor 192.0.2.7;  
    }  
}
```

4. On Router PE3, create an internal BGP peer group and specify the route reflector IP address (192.0.2.7) as the neighbor. Enable BGP to carry Layer 2 VPLS NLRI messages for this peer group and enable the processing of VPN-IPv4 addresses by including the **unicast** statement at the **[edit protocols bgp group group-name family inet-vpn]** hierarchy level.

```
[edit protocols]  
bgp {  
    group RR {  
        type internal;  
        local-address 192.0.2.3;  
        family inet-vpn {  
            unicast;  
        }  
        family l2vpn {  
            signaling;  
        }  
        neighbor 192.0.2.7;  
    }  
}
```

```
}
```

5. For the Layer 3 VPN domain on Router PE3 and Router PE5, enable RSVP on all interfaces. Disable RSVP on the fxp0.0 interface.

```
[edit protocols]
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

6. On Router PE3 and Router PE5, create label-switched paths (LSPs) to the route reflector and the other PE routers. The following example shows the configuration on Router PE5.

```
[edit protocols]
mpls {
  label-switched-path to-RR {
    to 192.0.2.7;
  }
  label-switched-path to-PE2 {
    to 192.0.2.2;
  }
  label-switched-path to-PE3 {
    to 192.0.2.3;
  }
  label-switched-path to-PE4 {
    to 192.0.2.4;
  }
  label-switched-path to-PE1 {
    to 192.0.2.1;
  }
}
```

7. On Routers PE1, PE2, PE3, and PE5, configure the core interfaces with an IPv4 address and enable the MPLS address family. The following example shows the configuration of the xe-0/1/0 interface on Router PE2.

```
[edit]
interfaces {
  xe-0/1/0 {
    unit 0 {
      family inet {
        address 10.10.2.2/30;
      }
      family mpls;
    }
  }
}
```

8. On Router PE2 and Router PE3, configure LDP for the Layer 2 VPN MPLS signaling protocol for all interfaces. Disable LDP on the fxp0.0 interface. (RSVP can also be used.)

```
[edit protocols]
ldp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

9. On the route reflector, create an internal BGP peer group and specify the PE routers IP addresses as the neighbors.

```
[edit]
protocols {
  bgp {
    group RR {
      type internal;
      local-address 192.0.2.7;
      family inet {
        unicast;
      }
      family inet-vpn {
        unicast;
      }
      family l2vpn {
        signaling;
      }
      cluster 192.0.2.7;
      neighbor 192.0.2.1;
      neighbor 192.0.2.2;
      neighbor 192.0.2.4;
      neighbor 192.0.2.5;
      neighbor 192.0.2.3;
    }
  }
}
```

10. On the route reflector, configure MPLS LSPs towards Routers PE3 and PE5 to resolve the BGP next hops from inet.3 routing table.

```
[edit]
protocols {
  mpls {
    label-switched-path to-pe3 {
      to 192.0.2.3;
    }
    label-switched-path to-pe5 {
      to 192.0.2.5;
    }
    interface all;
  }
}
```



## Configuring the VPN Interfaces

### Step-by-Step Procedure

Router PE2 is one end of the Layer 2 VPN. Router PE3 is performing the Layer 2 VPN stitching between the Layer 2 VPN and the Layer 3 VPN. Router PE3 uses the logical tunnel interface (lt interface) configured with different logical interface units applied under two different Layer 2 VPN instances. The packet is looped though the lt interface configured on Router PE3. The configuration of Router PE5 contains the PE-CE interface.

1. On Router PE2, configure the ge-1/0/2 interface encapsulation. Include the encapsulation statement and specify the **ethernet-ccc** option (**vlan-ccc** encapsulation is also supported) at the **[edit interfaces ge-1/0/2]** hierarchy level. The encapsulation should be the same in a whole Layer 2 VPN domain (Routers PE2 and PE3). Also, configure interface lo0.
 

```
[edit]
interfaces {
  ge-1/0/2 {
    encapsulation ethernet-ccc;
    unit 0;
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.2/24;
      }
    }
  }
}
```
2. On Router PE2, configure the routing instance at the **[edit routing-instances]** hierarchy level. Also, configure the Layer 2 VPN protocol at the **[edit routing-instances routing-instances-name protocols]** hierarchy level. Configure the remote site ID as 3. Site ID 3 represents Router PE3 (Hub-PE). The Layer 2 VPN is using LDP as the signaling protocol. Be aware that in the following example, both the routing instance and the protocol are named **l2vpn**.

```
[edit]
routing-instances {
  l2vpn { # routing instance
    instance-type l2vpn;
    interface ge-1/0/2.0;
    route-distinguisher 65000:2;
    vrf-target target:65000:2;
    protocols {
      l2vpn { # protocol
        encapsulation-type ethernet;
        site CE2 {
          site-identifier 2;
          interface ge-1/0/2.0 {
            remote-site-id 3;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```

3. On Router PE5, configure the Gigabit Ethernet interface for the PE-CE link **ge-2/0/0** and configure the **lo0** interface.

```

[edit interfaces]
ge-2/0/0 {
  unit 0 {
    family inet {
      address 198.51.100.8/24;
    }
  }
}
lo0 {
  unit 0 {
  }
}

```

4. On Router PE5, configure the Layer 3 VPN routing instance (**L3VPN**) at the **[edit routing-instances]** hierarchy level. Also configure BGP at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```

[edit]
routing-instances {
  L3VPN {
    instance-type vrf;
    interface ge-2/0/0.0;
    route-distinguisher 65000:5;
    vrf-target target:65000:2;
    vrf-table-label;
    protocols {
      bgp {
        group ce5 {
          neighbor 198.51.100.2 {
            peer-as 200;
          }
        }
      }
    }
  }
}

```

5. In an MX Series router, such as Router PE3, you must create the tunnel services interface to be used for tunnel services. To create the tunnel service interface, include the **bandwidth** statement and specify the amount of bandwidth to reserve for tunnel services in gigabits per second at the **[edit chassis fpc slot-number pic slot-number tunnel-services]** hierarchy level.

```

[edit]
chassis {
  dump-on-panic;
  fpc 1 {

```

```

pic 1 {
  tunnel-services {
    bandwidth 1g;
  }
}
}

```

6. On Router PE3, configure the Gigabit Ethernet interface.

Include the **address** statement at the **[edit interfaces ge-1/0/1.0 family inet]** hierarchy level and specify **198.51.100.9/24** as the IP address.

```

[edit]
interfaces {
  ge-1/0/1 {
    unit 0 {
      family inet {
        address 198.51.100.9/24;
      }
    }
  }
}

```

7. On Router PE3, configure the **lt-1/1/10.0** logical tunnel interface at the **[edit interfaces lt-1/1/10 unit 0]** hierarchy level. Router PE3 is the router that is *stitching* the Layer 2 VPN to the Layer 3 VPN using the logical tunnel interface. The configuration of the peer unit interfaces is what makes the interconnection.

To configure the interface, include the **encapsulation** statement and specify the **ethernet-ccc** option. Include the **peer-unit** statement and specify the logical interface unit 1 as the peer tunnel interface. Include the **family** statement and specify the **ccc** option.

```

[edit]
interfaces {
  lt-1/1/10 {
    unit 0 {
      encapsulation ethernet-ccc;
      peer-unit 1;
      family ccc;
    }
  }
}

```

8. On Router PE3, configure the **lt-1/1/10.1** logical tunnel interface at the **[edit interfaces lt-1/1/10 unit 1]** hierarchy level.

To configure the interface, include the **encapsulation** statement and specify the **ethernet** option. Include the **peer-unit** statement and specify the logical interface unit 0 as the peer tunnel interface. Include the **family** statement and specify the **inet** option. Include the **address** statement at the **[edit interfaces lt-1/1/10 unit 0]** hierarchy level and specify **198.51.100.7/24** as the IPv4 address.

```
[edit]
interfaces {
  lt-1/1/10 {
    unit 1 {
      encapsulation ethernet;
      peer-unit 0;
      family inet {
        address 198.51.100.7/24;
      }
    }
  }
}
```

9. On Router PE3, add the **lt** interface unit 1 to the routing instance at the **[edit routing-instances L3VPN]** hierarchy level. Configure the instance type as **vrf** with **lt** peer-unit 1 as a PE-CE interface to terminate the Layer 2 VPN on Router PE2 into the Layer 3 VPN on Router PE3.

```
[edit]
routing-instances {
  L3VPN {
    instance-type vrf;
    interface ge-1/0/1.0;
    interface lt-1/1/10.1;
    route-distinguisher 65000:33;
    vrf-target target:65000:2;
    vrf-table-label;
    protocols {
      bgp {
        export direct;
        group ce3 {
          neighbor 198.51.100.10 {
            peer-as 100;
          }
        }
      }
    }
  }
}
```

10. On Router PE3, add the **lt** interface unit 0 to the routing instance at the **[edit routing-instances protocols l2vpn]** hierarchy level. Also configure the same vrf target for the Layer 2 VPN and Layer 3 VPN routing instances, so that the routes can be leaked between the instances. The example configuration in the previous step shows the vrf target for the **L3VPN** routing instance. The following example shows the vrf target for the **l2vpn** routing instance.

```
[edit]
routing-instances {
  l2vpn {
    instance-type l2vpn;
    interface lt-1/1/10.0;
    route-distinguisher 65000:3;
    vrf-target target:65000:2;
```

```

protocols {
  l2vpn {
    encapsulation-type ethernet;
    site CE3 {
      site-identifier 3;
      interface lt-1/1/10.0 {
        remote-site-id 2;
      }
    }
  }
}

```

11. On Router PE3, configure the **policy-statement** statement to export the routes learned from the directly connected **lt** interface unit 1 to all the CE routers for connectivity, if needed.

```

[edit]
policy-options {
  policy-statement direct {
    term 1 {
      from protocol direct;
      then accept;
    }
  }
}

```

**Results** The following output shows the full configuration of Router PE2:

```

Router PE2 interfaces {
  xe-0/1/0 {
    unit 0 {
      family inet {
        address 10.10.2.2/30;
      }
      family mpls;
    }
  }
  xe-0/2/0 {
    unit 0 {
      family inet {
        address 10.10.5.1/30;
      }
      family mpls;
    }
  }
  xe-0/3/0 {
    unit 0 {
      family inet {
        address 10.10.4.1/30;
      }
      family mpls;
    }
  }
}

```

```
    }
  }
  ge-1/0/2 {
    encapsulation ethernet-ccc;
    unit 0;
  }
  fxp0 {
    apply-groups [ re0 re1 ];
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.2/24;
      }
    }
  }
}
routing-options {
  static {
    route 172.0.0.0/8 next-hop 172.19.59.1;
  }
  autonomous-system 65000;
}
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 192.0.2.2;
      family l2vpn {
        signaling;
      }
      neighbor 192.0.2.7;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

```

routing-instances {
  l2vpn {
    instance-type l2vpn;
    interface ge-1/0/2.0;
    route-distinguisher 65000:2;
    vrf-target target:65000:2;
    protocols {
      l2vpn {
        encapsulation-type ethernet;
        site CE2 {
          site-identifier 2;
          interface ge-1/0/2.0 {
            remote-site-id 3;
          }
        }
      }
    }
  }
}

```

The following output shows the final configuration of Router PE5:

```

Router PE5  interfaces {
              ge-0/0/0 {
                unit 0 {
                  family inet {
                    address 10.10.4.2/30;
                  }
                  family mpls;
                }
              }
              xe-0/1/0 {
                unit 0 {
                  family inet {
                    address 10.10.6.2/30;
                  }
                  family mpls;
                }
              }
              ge-1/0/0 {
                unit 0 {
                  family inet {
                    address 10.10.9.1/30;
                  }
                  family mpls;
                }
              }
              xe-1/1/0 {
                unit 0 {
                  family inet {
                    address 10.10.3.2/30;
                  }
                  family mpls;
                }
              }
              ge-2/0/0 {

```

```
    unit 0 {
      family inet {
        address 198.51.100.8/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.5/24;
      }
    }
  }
}
routing-options {
  static {
    route 172.0.0.0/8 next-hop 172.19.59.1;
  }
  autonomous-system 65000;
}
protocols {
  rsvp {
    interface all {
      link-protection;
    }
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-RR {
      to 192.0.2.7;
    }
    label-switched-path to-PE2 {
      to 192.0.2.2;
    }
    label-switched-path to-PE3 {
      to 192.0.2.3;
    }
    label-switched-path to-PE4 {
      to 192.0.2.4;
    }
    label-switched-path to-PE1 {
      to 192.0.2.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group to-rr {
      type internal;
      local-address 192.0.2.5;
      family inet-vpn {
        unicast;
      }
    }
  }
}
```



```

    }
    family l2vpn {
        signaling;
    }
    neighbor 192.0.2.7;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}
routing-instances {
    L3VPN {
        instance-type vrf;
        interface ge-2/0/0.0;
        route-distinguisher 65000:5;
        vrf-target target:65000:2;
        vrf-table-label;
        protocols {
            bgp {
                group ce5 {
                    neighbor 198.51.100.2 {
                        peer-as 200;
                    }
                }
            }
        }
    }
}
}
}

```

The following output shows the final configuration of Router PE3:

```

Router PE3  chassis {
              dump-on-panic;
              fpc 1 {
                  pic 1 {
                      tunnel-services {
                          bandwidth 1g;
                      }
                  }
              }
              network-services ip;
          }
          interfaces {

```

```
ge-1/0/1 {
  unit 0 {
    family inet {
      address 198.51.100.9/24;
    }
  }
}
lt-1/1/10 {
  unit 0 {
    encapsulation ethernet-ccc;
    peer-unit 1;
    family ccc;
  }
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet {
      address 198.51.100.7/24;
    }
  }
}
xe-2/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
    family mpls;
  }
}
xe-2/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.1/30;
    }
    family mpls;
  }
}
xe-2/2/0 {
  unit 0 {
    family inet {
      address 10.10.5.2/30;
    }
    family mpls;
  }
}
xe-2/3/0 {
  unit 0 {
    family inet {
      address 10.10.1.2/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
```

```

        address 192.0.2.3/24;
    }
}
}
routing-options {
    static {
        route 172.0.0.0/8 next-hop 172.19.59.1;
    }
    autonomous-system 65000;
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path to-RR {
            to 192.0.2.7;
        }
        label-switched-path to-PE2 {
            to 192.0.2.2;
        }
        label-switched-path to-PE5 {
            to 192.0.2.5;
        }
        label-switched-path to-PE4 {
            to 192.0.2.4;
        }
        label-switched-path to-PE1 {
            to 192.0.2.1;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group RR {
            type internal;
            local-address 192.0.2.3;
            family inet-vpn {
                unicast;
            }
            family l2vpn {
                signaling;
            }
            neighbor 192.0.2.7;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;

```

```
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}
policy-options {
    policy-statement direct {
        term 1 {
            from protocol direct;
            then accept;
        }
    }
}
routing-instances {
    L3VPN {
        instance-type vrf;
        interface ge-1/0/1.0;
        interface lt-1/1/10.1;
        route-distinguisher 65000:33;
        vrf-target target:65000:2;
        vrf-table-label;
        protocols {
            bgp {
                export direct;
                group ce3 {
                    neighbor 198.51.100.10 {
                        peer-as 100;
                    }
                }
            }
        }
    }
}
l2vpn {
    instance-type l2vpn;
    interface lt-1/1/10.0;
    route-distinguisher 65000:3;
    vrf-target target:65000:2;
    protocols {
        l2vpn {
            encapsulation-type ethernet;
            site CE3 {
                site-identifier 3;
                interface lt-1/1/10.0 {
                    remote-site-id 2;
                }
            }
        }
    }
}
}
```

```
}
```

## Verification

Verify the Layer 2 VPN-to-Layer 3 VPN interconnection:

- [Verifying Router PE2 VPN Interface on page 977](#)
- [Verifying Router PE3 VPN Interface on page 978](#)
- [Verifying End-to-End connectivity from Router CE2 to Router CE5 and Router CE3 on page 980](#)

### Verifying Router PE2 VPN Interface

**Purpose** Check that the Layer 2 VPN is up and working at the Router PE2 interface and that all the routes are there.

- Action** 1. Use the **show l2vpn connections** command to verify that the connection site ID is 3 for Router PE3 and that the status is **Up**.

```
user@PE2> show l2vpn connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch    WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down  NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch            MI -- Mesh-Group ID not available
BK -- Backup connection       ST -- Standby connection
PF -- Profile parse failure    PB -- Profile busy
RS -- remote site standby
```

Legend for interface status

```
Up -- operational
Dn -- down
```

Instance: l2vpn

Local site: CE2 (2)

connection-site	Type	St	Time last up	# Up trans
3	rmt	Up	Jan 7 14:14:37 2010	1

Remote PE: 192.0.2.3, Negotiated control-word: Yes (Null)  
Incoming label: 800000, Outgoing label: 800001  
Local interface: ge-1/0/2.0, Status: Up, Encapsulation: ETHERNET

2. Use the **show route table** command to verify that the Layer 2 VPN route is present and that there is a next hop of **10.10.5.2** through the **xe-0/2/0.0** interface. The following output verifies that the Layer 2 VPN routes are present in the **l2vpn.l2vpn.0** table. Similar output should be displayed for Router PE3.

```

user@PE2> show route table l2vpn.l2vpn.0
l2vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

65000:2:2:3/96
* [L2VPN/170/-101] 02:40:35, metric2 1
  Indirect
65000:3:3:1/96
* [BGP/170] 02:40:35, localpref 100, from 192.0.2.7
  AS path: I
  > to 10.10.5.2 via xe-0/2/0.0

```

3. Verify that Router PE2 has a Layer 2 VPN MPLS label pointing to the LDP label to Router PE3 in both directions (PUSH and POP).

```

user@PE2> show route table mpls.0
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0
* [MPLS/0] 1w3d 08:57:41, metric 1
  Receive
1
* [MPLS/0] 1w3d 08:57:41, metric 1
  Receive
2
* [MPLS/0] 1w3d 08:57:41, metric 1
  Receive
300560
* [LDP/9] 19:45:53, metric 1
  > to 10.10.2.1 via xe-0/1/0.0, Pop
300560(S=0)
* [LDP/9] 19:45:53, metric 1
  > to 10.10.2.1 via xe-0/1/0.0, Pop
301008
* [LDP/9] 19:45:53, metric 1
  > to 10.10.4.2 via xe-0/3/0.0, Swap 299856
301536
* [LDP/9] 19:45:53, metric 1
  > to 10.10.4.2 via xe-0/3/0.0, Pop
301536(S=0)
* [LDP/9] 19:45:53, metric 1
  > to 10.10.4.2 via xe-0/3/0.0, Pop
301712
* [LDP/9] 16:14:52, metric 1
  > to 10.10.5.2 via xe-0/2/0.0, Swap 315184
301728
* [LDP/9] 16:14:52, metric 1
  > to 10.10.5.2 via xe-0/2/0.0, Pop
301728(S=0)
* [LDP/9] 16:14:52, metric 1
  > to 10.10.5.2 via xe-0/2/0.0, Pop
800000
* [L2VPN/7] 02:40:35
  > via ge-1/0/2.0, Pop Offset: 4
ge-1/0/2.0
* [L2VPN/7] 02:40:35, metric2 1
  > to 10.10.5.2 via xe-0/2/0.0, Push 800001 Offset: -4

```

**Meaning** The `l2vpn` routing instance is up at interface `ge-1/0/2` and the Layer 2 VPN route is shown in table `l2vpn.l2vpn.0`. Table `mpls.0` shows the Layer 2 VPN routes used to forward the traffic using an LDP label.

### Verifying Router PE3 VPN Interface

**Purpose** Check that the Layer 2 VPN connection from Router PE2 and Router PE3 is **Up** and working.

- Action** 1. Verify that the BGP session with the route reflector for the family **l2vpn-signaling** and the family **inet-vpn** is established.

```
user@PE3> show bgp summary
```

```
Groups: 2 Peers: 2 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
bgp.l2vpn.0	1	1	0	0	0	0	0
bgp.L3VPN.0	1	1	0	0	0	0	0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn	State	#Active	/Received/Accepted/Damped...
192.0.2.7	65000	2063	2084	0	1	15:35:16	Establ		

```

  bgp.l2vpn.0: 1/1/1/0
  bgp.L3VPN.0: 1/1/1/0
  L3VPN.inet.0: 1/1/1/0
  l2vpn.l2vpn.0: 1/1/1/0

```

2. The following output verifies the Layer 2 VPN route and the label associated with it.

```
user@PE3> show route table l2vpn.l2vpn.0 detail
```

```
l2vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
65000:2:2:3/96 (1 entry, 1 announced)
```

```

  *BGP      Preference: 170/-101
             Route Distinguisher: 65000:2
             Next hop type: Indirect
             Next-hop reference count: 4
             Source: 192.0.2.7
             Protocol next hop: 192.0.2.2
             Indirect next hop: 2 no-forward
             State: <Secondary Active Int Ext>
             Local AS: 65000 Peer AS: 65000
             Age: 2:45:52 Metric2: 1
             Task: BGP_65000.192.0.2.7+60585
             Announcement bits (1): 0-l2vpn-l2vpn
             AS path: I (Originator) Cluster list: 192.0.2.7
             AS path: Originator ID: 192.0.2.2
             Communities: target:65000:2 Layer2-info: encaps:ETHERNET,
control flags:Control-Word, mtu: 0, site preference: 100 Accepted
             Label-base: 800000, range: 2, status-vector: 0x0
             Localpref: 100
             Router ID: 192.0.2.7
             Primary Routing Table bgp.l2vpn.0

```

3. The following output show the L2VPN MPLS.0 route in the mpls.0 route table.

```
user@PE3> show route table mpls.0
```

```
mpls.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```

0          *[MPLS/0] 1w3d 09:05:41, metric 1
             Receive
1          *[MPLS/0] 1w3d 09:05:41, metric 1
             Receive
2          *[MPLS/0] 1w3d 09:05:41, metric 1
             Receive
16         *[VPN/0] 15:59:24
             to table L3VPN.inet.0, Pop
315184     *[LDP/9] 16:21:53, metric 1
             > to 10.10.20.1 via xe-2/0/0.0, Pop
315184(S=0) *[LDP/9] 16:21:53, metric 1
             > to 10.10.20.1 via xe-2/0/0.0, Pop
315200     *[LDP/9] 01:13:44, metric 1

```

```

to 10.10.20.1 via xe-2/0/0.0, Swap 625297
> to 10.10.6.2 via xe-2/1/0.0, Swap 299856
315216 * [LDP/9] 16:21:53, metric 1
> to 10.10.6.2 via xe-2/1/0.0, Pop
315216(S=0) * [LDP/9] 16:21:53, metric 1
> to 10.10.6.2 via xe-2/1/0.0, Pop
315232 * [LDP/9] 16:21:45, metric 1
> to 10.10.1.1 via xe-2/3/0.0, Pop
315232(S=0) * [LDP/9] 16:21:45, metric 1
> to 10.10.1.1 via xe-2/3/0.0, Pop
315248 * [LDP/9] 16:21:53, metric 1
> to 10.10.5.1 via xe-2/2/0.0, Pop
315248(S=0) * [LDP/9] 16:21:53, metric 1
> to 10.10.5.1 via xe-2/2/0.0, Pop
315312 * [RSVP/7] 15:02:40, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315312(S=0) * [RSVP/7] 15:02:40, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315328 * [RSVP/7] 15:02:40, metric 1
> to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
315360 * [RSVP/7] 15:02:40, metric 1
> to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
316272 * [RSVP/7] 01:13:27, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
316272(S=0) * [RSVP/7] 01:13:27, metric 1
> to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
800001 * [L2VPN/7] 02:47:33
> via lt-1/1/10.0, Pop Offset: 4
lt-1/1/10.0 * [L2VPN/7] 02:47:33, metric2 1
> to 10.10.5.1 via xe-2/2/0.0, Push 800000 Offset: -4

```

4. Use the **show route table mpls.0** command with the **detail** option to see the BGP attributes of the route such as next-hop type and label operations.

```

user@PE5> show route table mpls.0 detail
lt-1/1/10.0 (1 entry, 1 announced)
  *L2VPN Preference: 7
    Next hop type: Indirect
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 607
    Next hop: 10.10.5.1 via xe-2/2/0.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 192.0.2.2
    Push 800000 Offset: -4
    Indirect next hop: 8cae0a0 1048574
    State: <Active Int>
    Age: 2:46:34 Metric2: 1
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC
    AS path: I
    Communities: target:65000:2 Layer2-info: encaps:ETHERNET,
control flags:Control-Word, mtu: 0, site preference: 100

```

### Verifying End-to-End connectivity from Router CE2 to Router CE5 and Router CE3

**Purpose** Check the connectivity between Routers CE2, CE3, and CE5.



- Action** 1. Ping the Router CE3 IP address from Router CE2.

```
user@CE2> ping 198.51.100.10 # CE3 IP address
PING 198.51.100.10 (198.51.100.10): 56 data bytes
64 bytes from 198.51.100.10: icmp_seq=0 ttl=63 time=0.708 ms
64 bytes from 198.51.100.10: icmp_seq=1 ttl=63 time=0.610 ms
```

2. Ping the Router CE5 IP address from Router CE2.

```
user@CE2> ping 198.51.100.2 # CE5 IP address
PING 198.51.100.2 (198.51.100.2): 56 data bytes
64 bytes from 198.51.100.2: icmp_seq=0 ttl=62 time=0.995 ms
64 bytes from 198.51.100.2: icmp_seq=1 ttl=62 time=1.005 ms
```

- Related Documentation**
- *Understanding Layer 2 VPNs*
  - [Understanding Layer 3 VPNs on page 123](#)
  - [Interconnecting Layer 2 VPNs with Layer 3 VPNs Overview on page 957](#)



## PART 10

# Troubleshooting Layer 3 VPNs

- [Troubleshooting Layer 3 VPNs on page 985](#)



# Troubleshooting Layer 3 VPNs

- [Diagnosing Common Layer 3 VPN Problems on page 985](#)
- [Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing on page 989](#)
- [Example: Troubleshooting Layer 3 VPNs on page 996](#)

## Diagnosing Common Layer 3 VPN Problems

---

**Problem**    **Description:** To troubleshoot problems in the Layer 3 VPN configuration, start at one end of the VPN (the local customer edge [CE] router) and follow the routes to the other end of the VPN (the remote CE router).

**Solution**    The following troubleshooting steps should help you diagnose common problems:

1. If you configured a routing protocol between the local provider edge (PE) and CE routers, check that the peering and adjacency are fully operational. When you do this, be sure to specify the name of the routing instance. For example, to check OSPF adjacencies, enter the **show ospf neighbor instance *routing-instance-name*** command on the PE router.

If the peering and adjacency are not fully operational, check the routing protocol configuration on the CE router and check the routing protocol configuration for the associated VPN routing instance on the PE router.

2. Check that the local CE and PE routers can ping each other.

To check that the local CE router can ping the VPN interface on the local PE router, use a **ping** command in the following format, specifying the IP address or name of the PE router:

```
user@host> ping (ip-address | host-name)
```

To check that the local PE router can ping the CE router, use a **ping** command in the following format, specifying the IP address or name of the CE router, the name of the interface used for the VPN, and the source IP address (the local address) in outgoing echo request packets:

```
user@host> ping ip-address interface interface local echo-address
```

Often, the peering or adjacency between the local CE and local PE routers must come up before a **ping** command is successful. To check that a link is operational in a lab setting, remove the interface from the VPN routing and forwarding (VRF) by deleting the **interface** statement from the **[edit routing-instance routing-instance-name]** hierarchy level and recommitting the configuration. Doing this removes the interface from the VPN. Then try the **ping** command again. If the command is successful, configure the interface back into the VPN and check the routing protocol configuration on the local CE and PE routers again.

3. On the local PE router, check that the routes from the local CE router are in the VRF table (*routing-instance-name.inet.0*):

```
user@host> show route table routing-instance-name.inet.0 <detail>
```

The following example shows the routing table entries. Here, the loopback address of the CE router is **10.255.14.155/32** and the routing protocol between the PE and CE routers is BGP. The entry looks like any ordinary BGP announcement.

```
10.255.14.155/32 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Nexthop: 192.168.197.141 via fe-1/0/0.0, selected
            State: <Active Ext>
            Peer AS:      1
            Age: 45:46
            Task: BGP_1.192.168.197.141+179
            Announcement bits (2): 0-BGP.0.0.0.0+179 1-KRT
            AS path: 1 I
            Localpref: 100
            Router ID: 10.255.14.155
```

If the routes from the local CE router are not present in the VRF routing table, check that the CE router is advertising routes to the PE router. If static routing is used between the CE and PE routers, make sure the proper static routes are configured.

4. On a remote PE router, check that the routes from the local CE router are present in the *bgp.l3vpn.0* routing table:

```
user@host> show route table bgp.l3vpn.0 extensive
10.255.14.175:3:10.255.14.155/32 (1 entry, 0 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 10.255.14.175:3
            Source: 10.255.14.175
            Nexthop: 192.168.192.1 via fe-1/1/2.0, selected
            Label-switched-path vpn07-vpn05
            Push 100004, Push 100005(top)
            State: <Active Int Ext>
            Local AS:      69 Peer AS:      69
            Age: 15:27      Metric2: 338
            Task: BGP_69.10.255.14.175+179
            AS path: 1 I
            Communities: target:69:100
            BGP next hop: 10.255.14.175
            Localpref: 100
            Router ID: 10.255.14.175
            Secondary tables: VPN-A.inet.0
```

The output of the **show route table bgp.l3vpn.0 extensive** command contains the following information specific to the VPN:

- In the prefix name (the first line of the output), the route distinguisher is added to the route prefix of the local CE router. Because the route distinguisher is unique within the Internet, the concatenation of the route distinguisher and IP prefix provides unique VPN-IP version 4 (IPv4) routing entries.
- The **Route Distinguisher** field lists the route distinguisher separately from the VPN-IPv4 address.
- The **label-switched-path** field shows the name of the label-switched path (LSP) used to carry the VPN traffic.
- The **Push** field shows both labels being carried in the VPN-IPv4 packet. The first label is the inner label, which is the VPN label that was assigned by the PE router. The second label is the outer label, which is an RSVP label.
- The **Communities** field lists the target community.
- The **Secondary tables** field lists other routing tables on this router into which this route has been installed.

If routes from the local CE router are not present in the bgp.l3vpn.0 routing table on the remote PE router, do the following:

- Check the VRF import filter on the remote PE router, which is configured in the **vrf-import** statement. (On the local PE router, you check the VRF export filter, which is configured with the **vrf-export** statement.)
- Check that there is an operational LSP or an LDP path between the PE routers. To do this, check that the IBGP next-hop addresses are in the inet.3 table.
- Check that the IBGP session between the PE routers is established and configured properly.
- Check for “hidden” routes, which usually means that routes were not labeled properly. To do this, use the **show route table bgp.l3vpn.0 hidden** command.
- Check that the inner label matches the inner VPN label that is assigned by the local PE router. To do this, use the **show route table mpls** command.

The following example shows the output of this command on the remote PE router. Here, the inner label is **100004**.

```
...
Push 100004, Push 10005 (top)
```

The following example shows the output of this command on the local PE router, which shows that the inner label of **100004** matches the inner label on the remote PE router:

```
...
100004          *[VPN/7] 06:56:25, metric 1
> to 192.168.197.141 via fe-1/0/0.0, Pop
```

5. On the remote PE router, check that the routes from the local CE router are present in the VRF table (*routing-instance-name.inet.0*):

```

user@host> show route table routing-instance-name.inet.0 detail
10.255.14.155/32 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 10.255.14.175:3
            Source: 10.255.14.175
            Nexthop: 192.168.192.1 via fe-1/1/2.0, selected
            label-switched-path vpn07-vpn05
            Push 100004, Push 100005(top)
            State: <Secondary Active Int Ext>
            Local AS: 69 Peer AS: 69
            Age: 1:16:22 Metric2: 338
            Task: BGP_69.10.255.14.175+179
            Announcement bits (2): 1-KRT 2-VPN-A-RIP
            AS path: 1 I
            Communities: target:69:100
            BGP next hop: 10.255.14.175
            Localpref: 100
            Router ID: 10.255.14.175
            Primary Routing Table bgp.l3vpn.0

```

In this routing table, the route distinguisher is no longer prepended to the prefix. The last line, **Primary Routing Table**, lists the table from which this route was learned.

If the routes are not present in this routing table, but were present in the `bgp.l3vpn.0` routing table on the local CE router, the routes might have not passed the VRF import policy on the remote PE router.

If a VPN-IPv4 route matches no **vrf-import** policy, the route does not show up in the `bgp.l3vpn` table at all and hence is not present in the VRF table. If this occurs, it might indicate that on the PE router, you have configured another **vrf-import** statement on another VPN (with a common target), and the routes show up in the `bgp.l3vpn.0` table, but are imported into the wrong VPN.

6. On the remote CE router, check that the routes from the local CE router are present in the routing table (`inet.0`):

```

user@host> show route

```

If the routes are not present, check the routing protocol configuration between the remote PE and CE routers, and make sure that peers and adjacencies (or static routes) between the PE and CE routers are correct.

7. If you determine that routes originated from the local CE router are correct, check the routes originated from the remote CE router by repeating this procedure.



## Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing

This example shows how to disable TTL decrementing in a single VRF routing instance in a Layer 3 VPN scenario.

- [Requirements on page 989](#)
- [Overview on page 989](#)
- [Configuration on page 990](#)
- [Verification on page 995](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Network Interfaces Configuration Guide*.

### Overview

To diagnose networking problems related to VPNs, it can be useful to disable normal time-to-live (TTL) decrementing. The IP header includes a TTL field that serves as a hop counter. At every routed hop, the TTL is decremented by one; if the TTL reaches zero before the packet reaches its destination, the packet is discarded and (optionally) an ICMP TTL exceeded message is sent to the source. MPLS labels also have a TTL field. MPLS routers copy the TTL of an IP packet when it enters a label-switched path (LSP). An IP packet with a TTL of 27 receives an MPLS label with a TTL of 27. Junos OS decrements the MPLS TTL of an MPLS-encapsulated packet in place of the IP TTL, at every label-switched hop. Because the MPLS TTL is copied (or propagated) from the IP TTL, a traceroute lists every hop in the path, be it routed or label-switched. When the packet exits the LSP, the decremented MPLS TTL is propagated back into the IP TTL field.

By default, TTL propagation is enabled. The global **no-propagate-ttl** statement disables TTL propagation at the router level and affects all RSVP-signalled or LDP-signalled LSPs. When a router acts as an ingress router for an LSP and the router configuration includes the **no-propagate-ttl** statement, the router pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When a router acts as the penultimate router, it pops the MPLS header without propagating the MPLS TTL into the IP packet. Thus the IP packet TTL value is preserved, regardless of the hop count of the LSP.

Instead of configuring TTL propagation behavior at the router level, you can configure the behavior for the routes in a VRF routing instance. This example shows how to disable TTL propagation for the routes in a single VRF routing instance instead of at the global router level.

The per-VRF configuration takes precedence over the global router configuration. If you disable TTL propagation on the router and explicitly enable TTL propagation for a single VRF routing instance, TTL propagation is in effect for that routing instance. To explicitly

enable TTL propagation on a VRF routing instance, include the **vrf-propagate-ttl** statement in the routing instance.


When you change the TTL propagation behavior, old next hops for VRF routes are deleted from the inet.3 routing table and new next hops are added.

You need only configure the **vrf-propagate-ttl** or **no-vrf-propagate-ttl** statement on the ingress routers.

### Topology Diagram

Figure 90 on page 990 shows the topology used in this example. Router PE1 and Router PE2 have two VPNs---VPN-A and VPN-B. Devices CE1 and CE4 belong to VPN-A. Devices CE2 and CE5 belong to VPN-B. In this example, Router PE1 has TTL propagation disabled on VPN-A but not on VPN-B. Packets received by PE1 on the interface connected to CE1 have TTL propagation disabled. This example shows the configuration on Router PE1. You do not need to include the **no-vrf-propagate-ttl** statement on the egress router (PE2).

Figure 90: Disabling TTL Propagation for a Single VPN



### Configuration

#### CLI Quick Configuration

To quickly disable TTL propagation in a VRF routing instance, copy the following commands and paste the commands into the CLI.

```
[edit]
set interfaces lo0 unit 0 family inet address 10.255.179.45/32 primary
set protocols mpls interface all
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.179.45
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp neighbor 10.255.179.71
set protocols ospf area 0.0.0.0 interface fe-1/1/2.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp interface all
set policy-options policy-statement VPN-A-export term a from protocol ospf
set policy-options policy-statement VPN-A-export term a from interface ge-1/2/0.0
set policy-options policy-statement VPN-A-export term a then community add VPN-A
set policy-options policy-statement VPN-A-export term b then reject
set policy-options policy-statement VPN-A-import term a from protocol bgp
set policy-options policy-statement VPN-A-import term a from community VPN-A
```

```

set policy-options policy-statement VPN-A-import term a then accept
set policy-options policy-statement VPN-A-import term b then reject
set policy-options policy-statement VPN-B-export term a from protocol static
set policy-options policy-statement VPN-B-export term a then community add VPN-B
set policy-options policy-statement VPN-B-export term b then accept
set policy-options policy-statement VPN-B-export term b then reject
set policy-options policy-statement VPN-B-import term a from protocol bgp
set policy-options policy-statement VPN-B-import term a from community VPN-B
set policy-options policy-statement VPN-B-import term a then accept
set policy-options policy-statement VPN-B-import term b then reject
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept
set policy-options community VPN-A members target:1:100
set policy-options community VPN-B members target:1:200
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface ge-1/2/0.0
set routing-instances VPN-A route-distinguisher 10.255.179.45:100
set routing-instances VPN-A interface ge-1/2/0.0
set routing-instances VPN-A no-vrf-propagate-ttl
set routing-instances VPN-A vrf-import VPN-A-import
set routing-instances VPN-A vrf-export VPN-A-export
set routing-instances VPN-A protocols ospf export bgp-to-ospf
set routing-instances VPN-A protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set routing-instances VPN-B instance-type vrf
set routing-instances VPN-B interface so-0/1/0.0
set routing-instances VPN-B route-distinguisher 10.255.179.45:300
set routing-instances VPN-B vrf-import VPN-B-import
set routing-instances VPN-B vrf-export VPN-B-export
set routing-instances VPN-B routing-options static route 10.255.179.15/32 next-hop
    so-0/1/0.0
set routing-options autonomous-system 1

```

**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*.

To configure a flow map:

1. Configure the loopback interface.

```

[edit]
user@PE1# edit interfaces
[edit interfaces]
user@PE1# set lo0 unit 0 family inet address 10.255.179.45/32 primary
user@PE1# exit

```

2. Configure the routing protocols.

The internal BGP neighbor address is the loopback interface address of Router PE2 in [Figure 90 on page 990](#).

```

[edit]
user@PE1# edit protocols
[edit protocols]
user@PE1# set mpls interface all

```

```

user@PE1# set bgp group ibgp type internal
user@PE1# set bgp group ibgp local-address 10.255.179.45
user@PE1# set bgp group ibgp family inet-vpn unicast
user@PE1# set bgp group ibgp neighbor 10.255.179.71
user@PE1# set ospf area 0.0.0.0 interface fe-1/1/2.0
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface lo0.0
user@PE1# set ldp interface all
user@PE1# exit

```

3. Configure routing policies for VPN-A and VPN-B.

```

[edit]
user@PE1# edit policy-options
[edit policy-options]
user@PE1# set policy-statement VPN-A-export term a from protocol ospf
user@PE1# set policy-statement VPN-A-export term a from interface ge-1/2/0.0
user@PE1# set policy-statement VPN-A-export term a then community add VPN-A
user@PE1# set policy-statement VPN-A-export term a then accept
user@PE1# set policy-statement VPN-A-export term b then reject
user@PE1# set policy-statement VPN-A-import term a from protocol bgp
user@PE1# set policy-statement VPN-A-import term a from community VPN-A
user@PE1# set policy-statement VPN-A-import term a then accept
user@PE1# set policy-statement VPN-A-import term b then reject
user@PE1# set policy-statement VPN-B-export term a from protocol static
user@PE1# set policy-statement VPN-B-export term a then community add VPN-B
user@PE1# set policy-statement VPN-B-export term a then accept
user@PE1# set policy-statement VPN-B-export term b then reject
user@PE1# set policy-statement VPN-B-import term a from protocol bgp
user@PE1# set policy-statement VPN-B-import term a from community VPN-B
user@PE1# set policy-statement VPN-B-import term a then accept
user@PE1# set policy-statement VPN-B-import term b then reject
user@PE1# set policy-statement bgp-to-ospf from protocol bgp
user@PE1# set policy-statement bgp-to-ospf then accept
user@PE1# set community VPN-A members target:1:100
user@PE1# set community VPN-B members target:1:200
user@PE1# exit

```

4. Configure the VPN-A and VPN-B routing instances, including the **no-vrf-propagate-ttl** statement in VPN-A.

```

[edit]
user@PE1# edit routing-instances
[edit routing-instances]
user@PE1# set VPN-A instance-type vrf
user@PE1# set VPN-A interface ge-1/2/0.0
user@PE1# set VPN-A route-distinguisher 10.255.179.45:100
user@PE1# set VPN-A interface ge-1/2/0.0
user@PE1# set VPN-A no-vrf-propagate-ttl
user@PE1# set VPN-A vrf-import VPN-A-import
user@PE1# set VPN-A vrf-export VPN-A-export
user@PE1# set VPN-A protocols ospf export bgp-to-ospf
user@PE1# set VPN-A protocols ospf area 0.0.0.0 interface ge-1/2/0.0
user@PE1# set VPN-B instance-type vrf
user@PE1# set VPN-B interface so-0/1/0.0

```

```

user@PE1# set VPN-B route-distinguisher 10.255.179.45:300
user@PE1# set VPN-B vrf-import VPN-B-import
user@PE1# set VPN-B vrf-export VPN-B-export
user@PE1# set VPN-B routing-options static route 10.255.179.15/32 next-hop
so-0/1/0.0
user@PE1# exit

```

5. Define the local autonomous system.

```

[edit]
user@PE1# edit routing-options
[edit routing-options]
user@PE1# set autonomous-system 1
user@PE1# exit

```

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

```

[edit]
user@PE1# commit

```

## Results

Confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-instances**, and **show routing-options** commands.

```

user@PE1# show interfaces
lo0 {
  unit 0 {
    family inet {
      address 10.255.179.45/32 {
        primary;
      }
    }
  }
}

user@PE1# show policy-options
policy-statement VPN-A-export {
  term a {
    from {
      protocol ospf;
      interface ge-1/2/0.0;
    }
    then {
      community add VPN-A;
      accept;
    }
  }
  term b {
    then reject;
  }
}
policy-statement VPN-A-import {

```

```
term a {
  from {
    protocol bgp;
    community VPN-A;
  }
  then accept;
}
term b {
  then reject;
}
}
policy-statement VPN-B-export {
  term a {
    from protocol static;
    then {
      community add VPN-B;
      accept;
    }
  }
  term b {
    then reject;
  }
}
policy-statement VPN-B-import {
  term a {
    from {
      protocol bgp;
      community VPN-B;
    }
    then accept;
  }
  term b {
    then reject;
  }
}
policy-statement bgp-to-ospf {
  from protocol bgp;
  then accept;
}
community VPN-A members target:1:100;
community VPN-B members target:1:200;

user@PE1# show protocols
mpls {
  interface all;
}
bgp {
  group ibgp {
    type internal;
    local-address 10.255.179.45;
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.179.71;
  }
}
```

```

ospf {
  area 0.0.0.0 {
    interface fe-1/1/2.0;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0;
  }
}
ldp {
  interface all;
}

user@PE1# show routing-instances
VPN-A {
  instance-type vrf;
  interface ge-1/2/0.0;
  no-vrf-propagate-ttl;
  route-distinguisher 10.255.179.45:100;
  vrf-import VPN-A-import;
  vrf-export VPN-A-export;
  protocols {
    ospf {
      export bgp-to-ospf;
      area 0.0.0.0 {
        interface ge-1/2/0.0;
      }
    }
  }
}
VPN-B {
  instance-type vrf;
  interface so-0/1/0.0;
  route-distinguisher 10.255.179.45:300;
  vrf-import VPN-B-import;
  vrf-export VPN-B-export;
  routing-options {
    static {
      route 10.255.179.15/32 next-hop so-0/1/0.0;
    }
  }
}

user@PE1# show routing-options
autonomous-system 1;

```

## Verification

To verify the operation, run the following commands:

- See the **TTL Action** field in the output of the **show route extensive table VPN-A** command.
- See the **TTL Action** field in the output of the **show route extensive table VPN-B** command.
- On Device CE1, run the **traceroute** command to Device CE4's loopback address.
- On Device CE4, run the **traceroute** command to Device CE1's loopback address.

- Related Documentation**
- [Disabling Normal TTL Decrementing](#)

---

## Example: Troubleshooting Layer 3 VPNs

---

This example shows how to use the **ping** command to check the accessibility of various routers in a VPN topology, and how to use the **traceroute** command to check the path that packets travel between the VPN routers.

- [Requirements on page 996](#)
- [Overview on page 996](#)
- [Pinging the CE Router from Another CE Router on page 998](#)
- [Pinging the Remote PE and CE Routers from the Local CE Router on page 999](#)
- [Pinging a CE Router from a Multiaccess Interface on page 1000](#)
- [Pinging the Directly Connected PE Routers from the CE Routers on page 1002](#)
- [Pinging the Directly Connected CE Routers from the PE Routers on page 1003](#)
- [Pinging the Remote CE Router from the Local PE Router on page 1004](#)
- [Troubleshooting Inconsistently Advertised Routes from Gigabit Ethernet Interfaces on page 1005](#)

### Requirements

This example uses the following hardware and software components:

- M Series routers
- Junos OS Release 10.0R1 and later

### Overview

---


#### Topology

---

The topology shown in [Figure 91 on page 997](#) illustrates the network used in this example to demonstrate how to employ the ping and traceroute commands to test connectivity between the routers participating in a Layer 3 VPN.



Figure 91: Layer 3 VPN Topology for ping and traceroute Examples



1666

## Pinging the CE Router from Another CE Router

**Step-by-Step Procedure** The following describes how to use the ping and **traceroute** commands to troubleshoot Layer 3 VPN topologies. You can ping one CE router from the other by specifying the other CE router's loopback address as the IP address in the **ping** command. This **ping** command succeeds if the loopback addresses have been announced by the CE routers to their directly connected PE routers. The success of these **ping** commands also means that Router CE1 can ping any network devices beyond Router CE2, and vice versa. [Figure 91 on page 997](#) shows the topology referenced in the following steps:

1. Ping Router CE2 (VPN5) from Router CE1 (VPN4):

```
user@vpn4> ping 10.255.10.5 local 10.255.10.4 count 3
PING 10.255.10.5 (10.255.10.5): 56 data bytes
64 bytes from 10.255.10.5: icmp_seq=0 ttl=253 time=1.086 ms
64 bytes from 10.255.10.5: icmp_seq=1 ttl=253 time=0.998 ms
64 bytes from 10.255.10.5: icmp_seq=2 ttl=253 time=1.140 ms
--- 10.255.10.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.998/1.075/1.140/0.059 ms
```

2. To determine the path from Router CE1's loopback interface to Router CE2's loopback interface, use the **traceroute** command:

```
user@vpn4> traceroute 10.255.10.5 source 10.255.10.4
traceroute to 10.255.10.5 (10.255.10.5) from 10.255.10.4, 30 hops max, 40 byte
packets
 1 vpn1-fe-110.isp-core.net (192.168.192.1) 0.680 ms 0.491 ms 0.456 ms
 2 vpn2-t3-001.isp-core.net (192.168.192.110) 0.857 ms 0.766 ms 0.754 ms

MPLS Label=100005 CoS=0 TTL=1 S=1
 3 vpn5.isp-core.net (10.255.10.5) 0.825 ms 0.886 ms 0.732 ms
```

3. When you use the **traceroute** command to examine the path used by a Layer 3 VPN, the provider (P) routers in the service provider's network are not displayed. As shown above, the jump from Router VPN1 to Router VPN2 is displayed as a single hop. The P router (VPN3) shown in [Figure 91 on page 997](#) is not displayed.

4. Ping Router CE1 (VPN4) from Router CE2 (VPN5):

```
user@vpn5> ping 10.255.10.4 local 10.255.10.5 count 3
PING 10.255.10.4 (10.255.10.4): 56 data bytes
64 bytes from 10.255.10.4: icmp_seq=0 ttl=253 time=1.042 ms
64 bytes from 10.255.10.4: icmp_seq=1 ttl=253 time=0.998 ms
64 bytes from 10.255.10.4: icmp_seq=2 ttl=253 time=0.954 ms
--- 10.255.10.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.954/0.998/1.042/0.036 ms
```

5. To determine the path from Router CE2 to Router CE1, use the **traceroute** command:

```
user@vpn5> traceroute 10.255.10.4 source 10.255.10.5
traceroute to 10.255.10.4 (10.255.10.4) from 10.255.10.5, 30 hops max, 40 byte
packets
```

```

1  vpn-08-t3-003.isp-core.net (192.168.193.2)  0.686 ms  0.519 ms  0.548 ms

2  vpn1-so-100.isp-core.net (192.168.192.100)  0.918 ms  0.869 ms  0.859 ms

      MPLS Label=100021 CoS=0 TTL=1 S=1
3  vpn4.isp-core.net (10.255.10.4)  0.878 ms  0.760 ms  0.739 ms

```

## Pinging the Remote PE and CE Routers from the Local CE Router

**Step-by-Step Procedure** From the local CE router, you can ping the VPN interfaces on the remote PE and CE routers, which are point-to-point interfaces. [Figure 91 on page 997](#) shows the topology referenced in the following examples:

1. Ping router CE2 from router CE1.

```

user@vpn4> ping 192.168.193.5 local 10.255.10.4 count 3
PING 192.168.193.5 (192.168.193.5): 56 data bytes
64 bytes from 192.168.193.5: icmp_seq=0 ttl=253 time=1.040 ms
64 bytes from 192.168.193.5: icmp_seq=1 ttl=253 time=0.891 ms
64 bytes from 192.168.193.5: icmp_seq=2 ttl=253 time=0.944 ms
--- 192.168.193.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.891/0.958/1.040/0.062 ms

```

2. To determine the path from Router CE1's loopback interface to Router CE2's directly connected interface, use the **tracert** command:

```

user@vpn4> tracert 192.168.193.5 source 10.255.10.4
tracert to 192.168.193.5 (192.168.193.5) from 10.255.10.4, 30 hops max,
40 byte packets
 1  vpn1-fe-110.isp-core.net (192.168.192.1)  0.669 ms  0.508 ms  0.457 ms
 2  vpn2-t3-001.isp-core.net (192.168.192.110)  0.851 ms  0.769 ms  0.750 ms

      MPLS Label=100000 CoS=0 TTL=1 S=1
 3  vpn5-t3-003.isp-core.net (192.168.193.5)  0.829 ms  0.838 ms  0.731 ms

```

3. Ping Router PE2 (VPN2) from Router CE1 (VPN4). In this case, packets that originate at Router CE1 go to Router PE2, then to Router CE2, and back to Router PE2 before Router PE2 can respond to Internet Control Message Protocol (ICMP) requests. You can verify this by using the **tracert** command.

```

user@vpn4> ping 192.168.193.2 local 10.255.10.4 count 3
PING 192.168.193.2 (192.168.193.2): 56 data bytes
64 bytes from 192.168.193.2: icmp_seq=0 ttl=254 time=1.080 ms
64 bytes from 192.168.193.2: icmp_seq=1 ttl=254 time=0.967 ms
64 bytes from 192.168.193.2: icmp_seq=2 ttl=254 time=0.983 ms
--- 192.168.193.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.967/1.010/1.080/0.050 ms

```

4. To determine the path from Router CE1 to Router PE2, use the **tracert** command:

```

user@vpn4> tracert 192.168.193.2 source 10.255.10.4
tracert to 192.168.193.2 (192.168.193.2) from 10.255.10.4, 30 hops max,
40 byte packets
 1  vpn1-fe-110.isp-core.net (192.168.192.1)  0.690 ms  0.490 ms  0.458 ms

```

```

2  vpn2-t3-003.isp-core.net (192.168.193.2)  0.846 ms  0.768 ms  0.749 ms
   MPLS Label=100000 CoS=0 TTL=1 S=1
3  vpn5-t3-003.isp-core.net (192.168.193.5)  0.643 ms  0.703 ms  0.600 ms
4  vpn-08-t3-003.isp-core.net (192.168.193.2)  0.810 ms  0.739 ms  0.729 ms

```

## Pinging a CE Router from a Multiaccess Interface

- Step-by-Step Procedure** You cannot ping one CE router from the other if the VPN interface is a multiaccess interface, such as the **fe-1/1/2.0** interface on Router CE1. To ping Router CE1 from Router CE2, you must either include the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy level on Router PE1 or configure a static route on Router PE1 to the VPN interface of Router CE1. If you include the **vrf-table-label** statement to ping a router, you cannot configure a static route.
1. If you configure a static route on Router PE1 to the VPN interface of Router CE1, its next hop must point to Router CE1 (at the **[edit routing-instance routing-instance-name]** hierarchy level), and this route must be announced from Router PE1 to Router PE2 as shown in the following configuration:

```

[edit]
routing-instances {
  direct-multipoint {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 69:1;
    vrf-import direct-import;
    vrf-export direct-export;
    routing-options {
      static {
        route 192.168.192.4/32 next-hop 192.168.192.4;
      }
    }
  }
  protocols {
    bgp {
      group to-vpn4 {
        peer-as 1;
        neighbor 192.168.192.4;
      }
    }
  }
}
policy-options {
  policy-statement direct-export {
    term a {
      from protocol bgp;
      then {
        community add direct-comm;
        accept;
      }
    }
    term b {
      from {
        protocol static;
        route-filter 192.168.192.4/32 exact;
      }
    }
  }
}

```

```

    }
    then {
        community add direct-comm;
        accept;
    }
}
term d {
    then reject;
}
}
}
}
}

```

2. Now you can ping Router CE1 from Router CE2:

```

user@vpn5> ping 192.168.192.4 local 10.255.10.5 count 3
PING 192.168.192.4 (192.168.192.4): 56 data bytes
64 bytes from 192.168.192.4: icmp_seq=0 ttl=253 time=1.092 ms
64 bytes from 192.168.192.4: icmp_seq=1 ttl=253 time=1.019 ms
64 bytes from 192.168.192.4: icmp_seq=2 ttl=253 time=1.031 ms
--- 192.168.192.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.019/1.047/1.092/0.032 ms

```

3. To determine the path between these two interfaces, use the **traceroute** command:

```

user@vpn5> traceroute 192.168.192.4 source 10.255.10.5
traceroute to 192.168.192.4 (192.168.192.4) from 10.255.10.5, 30 hops max,
40 byte packets
 1  vpn-08-t3003.isp-core.net (192.168.193.2)  0.678 ms  0.549 ms  0.494 ms

 2  vpn1-so-100.isp-core.net (192.168.192.100)  0.873 ms  0.847 ms  0.844 ms

    MPLS Label=100021 CoS=0 TTL=1 S=1
 3  vpn4-fe-112.isp-core.net (192.168.192.4)  0.825 ms  0.743 ms  0.764 ms

```

## Pinging the Directly Connected PE Routers from the CE Routers

**Step-by-Step Procedure** From the loopback interfaces on the CE routers, you can ping the VPN interface on the directly connected PE router. [Figure 91 on page 997](#) shows the topology referenced in this procedure:

1. From the loopback interface on Router CE1 (VPN4), ping the VPN interface, **fe-1/1/0.0**, on Router PE1:

```
user@vpn4> ping 192.168.192.1 local 10.255.10.4 count 3
PING 192.168.192.1 (192.168.192.1): 56 data bytes
64 bytes from 192.168.192.1: icmp_seq=0 ttl=255 time=0.885 ms
64 bytes from 192.168.192.1: icmp_seq=1 ttl=255 time=0.757 ms
64 bytes from 192.168.192.1: icmp_seq=2 ttl=255 time=0.734 ms
--- 192.168.192.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.734/0.792/0.885/0.066 ms
```

2. From the loopback interface on Router CE2 (VPN5), ping the VPN interface, **t3-0/0/3.0**, on Router PE2:

```
user@vpn5> ping 192.168.193.2 local 10.255.10.5 count 3
PING 192.168.193.2 (192.168.193.2): 56 data bytes
64 bytes from 192.168.193.2: icmp_seq=0 ttl=255 time=0.998 ms
64 bytes from 192.168.193.2: icmp_seq=1 ttl=255 time=0.834 ms
64 bytes from 192.168.193.2: icmp_seq=2 ttl=255 time=0.819 ms
--- 192.168.193.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.819/0.884/0.998/0.081 ms
```

3. From the loopback interface on Router CE2 (VPN5), ping the VPN interface, **t3-0/0/3.0**, on Router PE2:

```
user@vpn5> ping 192.168.193.2 local 10.255.10.5 count 3
PING 192.168.193.2 (192.168.193.2): 56 data bytes
64 bytes from 192.168.193.2: icmp_seq=0 ttl=255 time=0.998 ms
64 bytes from 192.168.193.2: icmp_seq=1 ttl=255 time=0.834 ms
64 bytes from 192.168.193.2: icmp_seq=2 ttl=255 time=0.819 ms
--- 192.168.193.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.819/0.884/0.998/0.081 ms
```

4. To determine the path from the loopback interface on Router CE2 to the VPN interfaces on Router PE2, use the **traceroute** command:

```
user@vpn5> traceroute 192.168.193.2 source 10.255.10.5
traceroute to 192.168.193.2 (192.168.193.2) from 10.255.10.5, 30 hops max,
40 byte packets
 1  vpn-08-t3003.isp-core.net (192.168.193.2)  0.852 ms  0.670 ms  0.656 ms
```

## Pinging the Directly Connected CE Routers from the PE Routers

**Step-by-Step Procedure** From the VPN and loopback interfaces on the PE routers, you can ping the VPN interface on the directly connected CE router. [Figure 91 on page 997](#) shows the topology referenced in this procedure:

1. From the VPN interface on the PE router (router PE1), you can ping the VPN or loopback interface on the directly connected CE router (router CE1).

From the VPN interface on Router PE1 (VPN1), ping the VPN interface, **fe-1/1/0.0**, on Router CE1:

```
user@vpn1> ping 192.168.192.4 interface fe-1/1/0.0 local 192.168.192.1 count 3
PING 192.168.192.4 (192.168.192.4): 56 data bytes
64 bytes from 192.168.192.4: icmp_seq=0 ttl=255 time=0.866 ms
64 bytes from 192.168.192.4: icmp_seq=1 ttl=255 time=0.728 ms
64 bytes from 192.168.192.4: icmp_seq=2 ttl=255 time=0.753 ms
--- 192.168.192.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.728/0.782/0.866/0.060 ms
```

2. From the VPN interface on Router PE1 (VPN1), ping the loopback interface, **10.255.10.4**, on Router CE1:

```
user@vpn1> ping 10.255.10.4 interface fe-1/1/0.0 local 192.168.192.1 count 3
PING 10.255.10.4 (10.255.10.4): 56 data bytes
64 bytes from 10.255.10.4: icmp_seq=0 ttl=255 time=0.838 ms
64 bytes from 10.255.10.4: icmp_seq=1 ttl=255 time=0.760 ms
64 bytes from 10.255.10.4: icmp_seq=2 ttl=255 time=0.771 ms
--- 10.255.10.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.760/0.790/0.838/0.034 ms
```

3. To determine the path from the VPN interface on Router PE1 to the VPN and loopback interfaces on Router CE1, respectively, use the following **traceroute** commands:

```
user@vpn1> traceroute 10.255.10.4 interface fe-1/1/0.0 source 192.168.192.1
traceroute to 10.255.10.4 (10.255.10.4) from 192.168.192.1, 30 hops max,
40 byte packets
 1  vpn4.isp-core.net (10.255.10.4)  0.842 ms  0.659 ms  0.621 ms
user@vpn1> traceroute 192.168.192.4 interface fe-1/1/0.0 source 192.168.192.1

traceroute to 192.168.192.4 (192.168.192.4) from 192.168.192.1, 30 hops max,
40 byte packets
 1  vpn4-fe-112.isp-core.net (192.168.192.4)  0.810 ms  0.662 ms  0.640 ms
```

4. From the VPN interface on Router PE2 (VPN2), ping the VPN interface, **t3-0/0/3.0**, on Router CE2:

```
user@vpn2> ping 192.168.193.5 interface t3-0/0/3.0 local 192.168.193.2 count 3
PING 192.168.193.5 (192.168.193.5): 56 data bytes
64 bytes from 192.168.193.5: icmp_seq=0 ttl=255 time=0.852 ms
64 bytes from 192.168.193.5: icmp_seq=1 ttl=255 time=0.909 ms
64 bytes from 192.168.193.5: icmp_seq=2 ttl=255 time=0.793 ms
--- 192.168.193.5 ping statistics ---
```

```
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.793/0.851/0.909/0.047 ms
```

- From the VPN interface on Router PE2 (VPN2), ping the loopback interface, **10.255.10.5**, on Router CE2:

```
user@vpn2> ping 10.255.10.5 interface t3-0/0/3.0 local 192.168.193.2 count 3
PING 10.255.10.5 (10.255.10.5): 56 data bytes
64 bytes from 10.255.10.5: icmp_seq=0 ttl=255 time=0.914 ms
64 bytes from 10.255.10.5: icmp_seq=1 ttl=255 time=0.888 ms
64 bytes from 10.255.10.5: icmp_seq=2 ttl=255 time=1.066 ms
--- 10.255.10.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.888/0.956/1.066/0.079 ms
```

- To determine the path from the VPN interface on Router PE2 to the VPN and loopback interfaces on Router CE2, respectively, use the following **traceroute** commands:

```
user@vpn2> traceroute 10.255.10.5 interface t3-0/0/3.0 source 192.168.193.2
traceroute to 10.255.10.5 (10.255.10.5) from 192.168.193.2, 30 hops max,
40 byte packets
 1  vpn5.isp-core.net (10.255.10.5)  1.009 ms  0.677 ms  0.633 ms
user@vpn2> traceroute 192.168.193.5 interface t3-0/0/3.0 source 192.168.193.2

traceroute to 192.168.193.5 (192.168.193.5) from 192.168.193.2, 30 hops max,
40 byte packets
 1  vpn5-t3-003.isp-core.net (192.168.193.5)  0.974 ms  0.665 ms  0.619 ms
```

## Pinging the Remote CE Router from the Local PE Router

**Step-by-Step Procedure** The following procedure is effective for Layer 3 VPNs only. To ping a remote CE router from a local PE router in a Layer 3 VPN, you need to configure the following interfaces:

- Configure a logical unit for the loopback interface.

To configure an additional logical unit on the loopback interface of the PE router, configure the **unit** statement at the **[edit interfaces lo0]** hierarchy level:

```
[edit interfaces]
lo0 {
  unit number {
    family inet {
      address address;
    }
  }
}
```

- Configure the loopback interface for the Layer 3 VPN routing instance on the local PE router. You can associate one logical loopback interface with each Layer 3 VPN routing instance, enabling you to ping a specific routing instance on a router.

Specify the loopback interface you configured in Step 1 using the **interface** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level:



```
[edit routing-instances routing-instance-name]
  interface interface-name;
```

The *interface-name* is the logical unit on the loopback interface (for example, **lo0.1**).

3. From the VPN interface on PE router, you can now ping the logical unit on the loopback interface on the remote CE router:

```
user@host> ping interface interface host
```

Use *interface* to specify the new logical unit on the loopback interface (for example, **lo0.1**). For more information about how to use the **ping interface** command, see the *Junos Interfaces Command Reference*.

## Troubleshooting Inconsistently Advertised Routes from Gigabit Ethernet Interfaces

**Step-by-Step Procedure** For direct routes on a LAN in a Layer 3 VPN, the Junos OS attempts to locate a CE router that can be designated as the next hop. If this cannot be done, advertised routes from Gigabit Ethernet interfaces are dropped.

In such instances:

1. Use the **static** statement at the **[edit routing-options]** or **[edit logical-systems logical-system-name routing-options]** hierarchy levels in the VRF routing instance to a CE router on the LAN subnet, configuring the CE router as the next hop. All traffic to directly destinations on this LAN will go to the CE router. You can add two static routes to two CE routers on the LAN for redundancy.
2. Configure the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy levels to map the inner label of a packet to a specific VRF routing table. This allows the examination of the encapsulated IP header to force IP lookups on the VRF routing instance for all traffic.



**NOTE:** The **vrf-table-label** statement is not available for every core-facing interface; for example, channelized interfaces are not supported. See [“Support on Ethernet, SONET/SDH, and T1/T3/E3 Interfaces for IP-Based Filtering” on page 162](#) for information about support for the **vrf-table-label** statement over Ethernet and SONET/SDH interfaces.



## PART 11

# Configuration Statements and Operational Commands

- [Configuration Statements \(All VPNs\) on page 1009](#)
- [Configuration Statements \(Layer 3 VPNs\) on page 1037](#)
- [Operational Commands on page 1159](#)



## CHAPTER 49

# Configuration Statements (All VPNs)


- [aggregate-label](#) on page 1010
- [backup-neighbor](#) on page 1011
- [description](#) (Routing Instances) on page 1012
- [family route-target](#) on page 1013
- [graceful-restart](#) (Enabling Globally) on page 1015
- [instance-type](#) on page 1017
- [interface](#) (Routing Instances) on page 1019
- [no-forwarding](#) on page 1020
- [proxy-generate](#) on page 1021
- [revert-time](#) (Protocols Layer 2 Circuits) on page 1022
- [route-distinguisher](#) on page 1024
- [route-distinguisher-id](#) on page 1026
- [route-target-filter](#) on page 1027
- [switchover-delay](#) on page 1028
- [unicast-reverse-path](#) on page 1029
- [vpn-apply-export](#) on page 1030
- [vrf-export](#) on page 1031
- [vrf-import](#) on page 1032
- [vrf-mtu-check](#) on page 1033
- [vrf-target](#) on page 1034

## aggregate-label

---

<b>Syntax</b>	<pre>aggregate-label {     community <i>community-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn6 unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp family inet6 labeled-unicast], [edit protocols bgp family inet-vpn unicast], [edit protocols bgp family inet6-vpn unicast]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify matching criteria (in the form of a community) such that all routes which match are assigned the same VPN label, selected from one of the several routes in the set defined by this criteria. This reduces the number of VPN labels that the router must consider, and aggregates the received labels.
<b>Options</b>	<b>community <i>community-name</i></b> —Specify the name of the community to which to apply the aggregate label.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Aggregate Labels for VPNs on page 21</a></li></ul>

## backup-neighbor

<b>Syntax</b>	<pre> backup-neighbor <i>address</i> {     community <i>name</i>;     mtu <i>number</i>;     hot-standby;     psn-tunnel-endpoint <i>address</i>;     standby;     static;     virtual-circuit-id <i>number</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>],</p> <p>[edit protocols l2circuit local-switching interface <i>interface-name</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 12.3 at the <b>[edit protocols l2circuit local-switching interface <i>interface-name</i>]</b> hierarchy level.</p>
<b>Description</b>	<p>Configure pseudowire redundancy for Layer 2 circuits and VPLS. A redundant pseudowire can act as a backup connection and can be configured between a PE router and a CE device or between PE routers, maintaining Layer 2 circuit and VPLS services after certain types of failures. This feature can help improve the reliability of certain types of networks where a single point of failure could interrupt service for customers.</p>
	<p> <b>NOTE:</b> VPLS is not supported on ACX Series routers. The <b>psn-tunnel-endpoint</b> statement is not supported at the <b>[edit protocols l2circuit local-switching interface <i>interface-name</i> end-interface interface <i>interface-name</i>]</b> hierarchy level.</p>
<b>Options</b>	<p><b><i>address</i></b>—Specifies the address for the backup neighbor.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring Pseudowire Redundancy on the PE Router</li> </ul>

- *Example: Configuring Layer 2 Circuit Switching Protection*
- *community (Protocols Layer 2 Circuit)*
- *psn-tunnel-endpoint*
- *virtual-circuit-id*

---

## description (Routing Instances)

---

<b>Syntax</b>	description text;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Provide a text description for the routing instance. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the <b>show route instance detail</b> command and has no effect on the operation of the routing instance.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.




## family route-target

<b>Syntax</b>	<pre>family route-target {     advertise-default;     external-paths <i>number</i>;     prefix-limit <i>number</i>;     proxy-generate &lt;route-target-policy <i>route-target-policy-name</i>&gt;; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Enable BGP route target filtering on the VPN.</p> <p>The <b>family route-target</b> statement is useful for filtering VPN routes before they are sent. Provider edge (PE) routers inform the route reflector (RR) which routes to send, using <b>family route-target</b> to provide the route-target-interest information. The RR then sends to the PE router only the advertisements containing the specified route target.</p>
<b>Options</b>	<p><b>advertise-default</b>—Cause the router to advertise the default route target route (<b>0:0:0/0</b>) and suppress all routes that are more specific. This can be used by a route reflector on BGP groups consisting of neighbors that act as provider edge (PE) routers only. PE routers often need to advertise all routes to the route reflector. Suppressing all route target advertisements other than the default route reduces the amount of information exchanged between the route reflector and the PE routers. The Junos OS further helps to reduce route target advertisement overhead by not maintaining dependency information unless a nondefault route is received.</p> <p><b>external-paths <i>number</i></b>—Cause the router to advertise the VPN routes that reference a given route target. The number you specify with the <b>external-paths</b> statement determines the number of external peer routers (currently advertising that route target) that receive the VPN routes. The default value is 1.</p> <p><b>prefix-limit <i>number</i></b>—The number of prefixes that can be received from a peer router.</p> <p>The remaining statement is described separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BGP Route Target Filtering for VPNs on page 38</a></li> <li>• <a href="#">Understanding Proxy BGP Route Target Filtering for VPNs on page 50</a></li> </ul>

- [Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs on page 67](#)

## graceful-restart (Enabling Globally)

<b>Syntax</b>	<pre> graceful-restart {   disable;   helper-disable;   maximum-helper-recovery-time <i>seconds</i>;   maximum-helper-restart-time <i>seconds</i>;   notify-duration <i>seconds</i>;   recovery-time <i>seconds</i>;   restart-duration <i>seconds</i>;   stale-routes-time <i>seconds</i>; } </pre>
<b>Hierarchy Level</b>	<pre> [edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. Because all routing protocols are not usually run on every routing instance, you must also configure graceful restart for individual routing protocols running on a routing instance, including the main routing instance. You can, optionally, modify the global settings at the individual protocol level.</p>
	<div>  <p><b>NOTE:</b></p> <ul style="list-style-type: none"> <li>For VPNs, the <code>graceful-restart</code> statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</li> <li>For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</li> <li>LDP sessions flap when <code>graceful-restart</code> configurations change.</li> </ul> </div>
<b>Default</b>	Graceful restart is disabled by default.
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Configuring Graceful Restart for QFabric Systems*

## instance-type

<b>Syntax</b>	<code>instance-type type;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>virtual-switch</b> and <b>layer2-control</b> options introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. <b>mpls-internet-multicast</b> option introduced in Junos OS Release 11.1 for the EX Series, M Series, MX Series, and T Series. <b>evpn</b> option introduced in Junos OS Release 13.2 for MX 3D Series routers. <b>evpn</b> option introduced in Junos OS Release 17.3 for the QFX Series. <b>forwarding</b> option introduced in Junos OS Release 14.2 for the PTX Series. <b>mpls-forwarding</b> option introduced in Junos OS Release 16.1 for the MX Series. <b>evpn-vpws</b> option introduced in Junos OS Release 17.1 for MX Series routers. Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.
<b>Description</b>	Define the type of routing instance.

### Options



**NOTE:** On ACX Series routers, you can configure only the forwarding, virtual router, and VRF routing instances.

**type**—Can be one of the following:

- **evpn**—(MX 3D Series routers, QFX switches and EX9200 switches)—Enable an Ethernet VPN (EVPN) on the routing instance.  
hierarchy level.
- **evpn-vpws**—Enable an Ethernet VPN (EVPN) Virtual Private Wire Service (VPWS) on the routing instance.
- **forwarding**—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- **l2backhaul-vpn**—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.

- **l2vpn**—Enable a Layer 2 VPN on the routing instance. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **layer2-control**—(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.
- **mpls-forwarding**—(MX Series routers only) Allow filtering and translation of route distinguisher (RD) values in IPv4 and IPv6 VPN address families on both routes received and routes sent for selected BGP sessions. In particular, for Inter-AS VPN Option-B networks, this option can prevent the malicious injection of VPN labels from one peer AS boundary router to another.
- **mpls-internet-multicast**—(EX Series, M Series, MX Series, and T Series routers only) Provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.
- **no-forwarding**—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.
- **virtual-router**—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the **interface** statement for this type of routing instance. You do not need to configure the **route-distinguisher**, **vrf-import**, and **vrf-export** statements.
- **virtual-switch**—(MX Series routers, EX9200 switches, and QFX switches only) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.
- **vpls**—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **vrf**—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (**instance-name.inet.0**) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Instance Type on page 10](#)
- [Configuring EVPN Routing Instances](#)
- [Configuring EVPN Routing Instances on EX9200 Switches](#)
- [Configuring Virtual Router Routing Instances](#)
- [Example: Configuring Filter-Based Forwarding on the Source Address](#)
- [Example: Configuring Filter-Based Forwarding on Logical Systems](#)
- [vpws-service-id](#)

## interface (Routing Instances)

**Syntax** `interface interface-name {  
    description text;  
}`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],  
[edit routing-instances *routing-instance-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.2 for EX Series switches.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.  
Statement introduced in Junos OS Release 13.2 for MX 3D Series routers.

**Description** Specify the interface over which the VPN traffic travels between the PE device and CE device. You configure the interface on the PE device. If the value **vrf** is specified for the **instance-type** statement included in the routing instance configuration, this statement is required.

**Options** *interface-name*—Name of the interface.

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Interfaces for VPN Routing on page 11](#)
- [Configuring EVPN Routing Instances](#)
- [Configuring EVPN Routing Instances on EX9200 Switches](#)
- [interface \(VPLS Routing Instances\)](#)

## no-forwarding

---

<b>Syntax</b>	no-forwarding;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Do not add ingress routes to the inet.0 routing table even if <b>traffic-engineering bgp-igp</b> (configured at the [edit protocols mpls] hierarchy level) is enabled.
<b>Default</b>	The <b>no-forwarding</b> statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when <b>traffic-engineering bgp-igp</b> is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Preventing Addition of Ingress Routes to the inet.0 Routing Table</a></li><li>• <a href="#">Configuring a Routing Protocol Between the Service Provider Routers on page 15</a></li></ul>



## proxy-generate

<b>Syntax</b>	<code>proxy-generate &lt;route-target-policy <i>route-target-policy-name</i>&gt;;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>family route-target</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family route-target</b>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>family route-target</b>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family route-target</b>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	Enable proxy BGP route target filtering (also known as proxy route target constrain, or proxy RTC). This feature is useful if you have a network environment where route target filtering is not widely deployed or fully supported. When configured for proxy BGP route target filtering, the device creates route target membership (RT membership) on behalf of its peers that do not have the route target filtering functionality. The device then distributes the RT membership advertisements from incoming BGP VPN routes to other devices in the network that need them.
<b>Options</b>	<b>route-target-policy <i>route-target-policy-name</i></b> —(Optional) Apply a routing policy that defines a subset of VPN routes to be used in your proxy BGP route target filter. The subset of VPN routes control which proxy BGP route targets are used to create the proxy BGP route target routes.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Proxy BGP Route Target Filtering for VPNs on page 51</a></li> <li>• <a href="#">Example: Configuring an Export Policy for BGP Route Target Filtering for VPNs on page 67</a></li> <li>• <a href="#">Configuring BGP Route Target Filtering for VPNs on page 37</a></li> <li>• <a href="#">family route-target on page 1013</a></li> <li>• <a href="#">rtf-prefix-list</a></li> </ul>

## revert-time (Protocols Layer 2 Circuits)

---

<b>Syntax</b>	<code>revert-time <i>seconds</i> maximum <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>],</code> <code>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2. <b>maximum</b> option introduced in Junos OS Release 13.2.
<b>Description</b>	<p>Specify a revert time for redundant Layer 2 circuits and VPLS pseudowires. When you have configured redundant pseudowires for Layer 2 circuits or VPLS, traffic is switched to the backup connection in the event that the primary connection fails. If you configure a revert time, when the configured time expires traffic is reverted to the primary path, assuming the primary path has been restored.</p> <p>With the <b>maximum</b> option, specify a maximum reversion interval to add after the <b>revert-time</b> delay. If a revert-time delay is defined but a maximum timer is not defined, VCs are restored upon the revert-timer's expiration.</p> <p>To reduce as much as possible the amount of traffic discarded, and potential data-path asymmetries observed during primary-to-backup transition periods, you can use this restoration timer. This restoration timer is activated when the backup path is performing as active, and then the primary path is restored. The goal is to avoid moving traffic back to the primary path right away, to make sure that the control plane's related tasks (such as IGP, LDP, RSVP, and internal BGP) have enough time to complete their updating cycle.</p> <p>By enabling a gradual return of traffic to the primary path, you can ensure that the relatively-slow control-plane processing and updating does not have a negative impact on the restoration process.</p> <p>The <b>maximum</b> option extends the revert timer's functionality to provide a jittered interval over which a certain number of circuits can be transitioned back to the primary path. By making use of this maximum value, you can define a time interval during which circuits are expected to switch over. As a consequence, circuits' effective transitions are scattered during restoration periods.</p> <p>When making use of <b>revert-time x maximum y</b> statement, you can ensure that the corresponding circuit that is active is moved to the primary path within a time-slot (t1) such as that: <math>x \leq t1 \leq y</math>. In other words, by activating this statement, you can ensure the following:</p> <ul style="list-style-type: none"><li>• VCs stay in the backup path for at least x seconds after the primary path comes back up.</li><li>• VCs are moved back to the primary path before y seconds have elapsed.</li></ul>

- $y \text{ maximum value} = x \text{ maximum value} * 2 = 1200 \text{ seconds}$ .

The ideal values for x and y will be conditioned to internal aspects of your network. For this reason, there are no default values for these settings. If no revert-time is set, the default behavior is non-revertive. That is, circuits are not returned to the primary path upon restoration. They are kept on the backup path.

**Default** Without the **revert-time** statement, virtual circuit (VC) traffic is not transitioned to the primary path upon restoration of the primary path.

**Options** *seconds*—Revert time in seconds.

**Range:** 0 through 600 seconds

maximum *seconds*—Number of seconds to delay path restoration after the **revert-time** delay.

**Range:** 0 through 1200 seconds

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS*

## route-distinguisher

<b>Syntax</b>	<code>route-distinguisher (as-number:id   ip-address:id);</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>],  [edit routing-instances <i>routing-instance-name</i>],  [edit routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>],  [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Support at [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>] hierarchy level introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Support at [edit routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>] hierarchy level introduced in Junos OS Release 13.2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.</p>
<b>Description</b>	<p>Specify an identifier attached to a route, enabling you to distinguish to which VPN or virtual private LAN service (VPLS) the route belongs. Each routing instance must have a unique route distinguisher (RD) associated with it. The RD is used to place bounds around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. If the instance type is <b>vrf</b>, the <b>route-distinguisher</b> statement is required.</p> <p>For Layer 2 VPNs and VPLS, if you configure the <b>l2vpn-use-bgp-rules</b> statement, you must configure a unique RD for each PE router participating in the routing instance.</p> <p>For other types of VPNs, we recommend that you use a unique RD for each provider edge (PE) router participating in specific routing instance. Although you can use the same RD on all PE routers for the same VPN routing instance, if you use a unique RD, you can determine the customer edge (CE) router from which a route originated within the VPN.</p> <p>For Layer 2 VPNs and VPLSs, if you configure mesh groups, the RD in each mesh group must be unique.</p>



**CAUTION:** We strongly recommend that if you change an RD that has already been configured, make the change during a maintenance window, as follows:

1. Deactivate the routing instance.
2. Change the RD.
3. Activate the routing instance.

This is not required if you are configuring the RD for the first time.

**Options** *as-number:number*—*as-number* is an assigned AS number, and *number* is any 2-byte or 4-byte value. The AS number can be from 1 through 4,294,967,295. If the AS number is a 2-byte value, the administrative number is a 4-byte value. If the AS number is a 4-byte value, the administrative number is a 2-byte value. An RD consisting of a 4-byte AS number and a 2-byte administrative number is defined as a type 2 RD in RFC 4364 *BGP/MPLS IP VPNs*.



**NOTE:** In Junos OS Release 9.1 and later, the numeric range for AS numbers is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. All releases of Junos OS support 2-byte AS numbers. To configure an RD that includes a 4-byte AS number, append the letter “L” to the end of the AS number. For example, an RD with the 4-byte AS number 7,765,000 and an administrative number of 1,000 is represented as 77765000L:1000.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in the plain-number format is represented as 1.10 in AS dot notation format.

*ip-address:id*—IP address (*ip-address* is a 4-byte value) within your assigned prefix range and a 2-byte value for the *id*. The IP address can be any globally unique unicast address.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ ). If the router you are configuring is a BGP peer of a router that does not support 4-byte AS numbers, you need to configure a local AS number. For more information, see *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring BGP Route Target Filtering for VPNs on page 41](#)
- [Example: Configuring FEC 129 BGP Autodiscovery for VPWS](#)
- [Configuring EVPN Routing Instances](#)
- [Configuring the Route Distinguisher on page 13](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\)](#)
- [Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\)](#)
- [l2vpn-use-bgp-rules](#)

## route-distinguisher-id

---

Syntax	route-distinguisher-id <i>ip-address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Automatically assign a route distinguisher to the routing instance.</p> <p>If you configure the <b>route-distinguisher</b> statement in addition to the <b>route-distinguisher-id</b> statement, the value configured for <b>route-distinguisher</b> supersedes the value generated from <b>route-distinguisher-id</b>.</p> <div><p><b>NOTE:</b> To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the <b>route-distinguisher</b> statement is different from the first half of the route distinguisher obtained by configuring the <b>route-distinguisher-id</b> statement.</p></div>
Options	<i>ip-address</i> —Address for routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BGP Route Target Filtering for VPNs on page 41</a></li><li>• <a href="#">Configuring the Route Distinguisher on page 13</a></li></ul>

## route-target-filter

<b>Syntax</b>	<pre>route-target-filter <i>destination</i> {     group <i>group-name</i>;     local;     neighbor <i>address</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options rib bgp.rtarget.0 static], [edit routing-options rib bgp.rtarget.0 static]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Statically configure route target filtering. Route target filtering allows you to distribute VPN routes to only the routers that need them. In VPN networks without route target filtering configured, BGP distributes all static VPN routes to all of the VPN peer routers. You can add static routes to the bgp.rtarget.0 routing table with specific NLRI-imposed constraints.
<b>Options</b>	<p><b><i>destination</i></b>—Allows you to specify the static route destination. This value must be in the format x:y/len. The x value is either an IP address or an AS number followed by an optional L to indicate a 4 byte AS number, and y is a number (for example, 123456L:100/64).</p> <p><b><i>group group-name(s)</i></b>—Installs an RT-Constrain filter for the destination for all peers in the specified BGP group. The route and corresponding BGP group are displayed in the output of the <b>show bgp group rtf detail</b> command.</p> <p><b><i>local</i></b>—Causes the router to originate the route target constrain NLRI, but does not install any filtering state for the prefix. This behavior can be useful when the router should always receive VPN routes with this route-target regardless of the state of a given BGP peering session or group status. The route is not displayed in the output of the <b>show bgp group rtf detail</b> command unless it is also included in either the BGP neighbor or BGP group configuration.</p> <p><b><i>neighbor address</i></b>—Installs an RT-Constrain filter for the destination for this BGP neighbor. The route is displayed in the output of the <b>show bgp group rtf detail</b> command. Specify the BGP neighbor using the router's IP address.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Static Route Target Filtering for VPNs on page 50</a></li> <li>• <a href="#">Reducing Network Resource Use with Static Route Target Filtering for VPNs on page 84</a></li> </ul>

## switchover-delay

---

<b>Syntax</b>	<code>switchover-delay <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> ], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	After the primary pseudowire goes down, specifies the delay (in milliseconds) to wait before the backup pseudowire takes over. You configure this statement for each backup neighbor configuration to adjust the switchover time after a failure is detected.
<b>Options</b>	<b><i>milliseconds</i></b> —Specify the time to wait before switching to the backup pseudowire after the primary pseudowire fails. <b>Default:</b> 10,000 milliseconds <b>Range:</b> 0 through 180,000 milliseconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Switchover Delay for the Pseudowires</i></li></ul>



## unicast-reverse-path

<b>Syntax</b>	unicast-reverse-path (active-paths   feasible-paths);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">forwarding-table</a> ], [edit routing-instances <i>routing-instance-name</i> instance-type <i>name</i> routing-options <a href="#">forwarding-table</a> ], [edit routing-options <a href="#">forwarding-table</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for routing instances added in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Control the operation of unicast reverse-path-forwarding check. This statement enables the RPF check to be used when routing is asymmetrical.
<b>Options</b>	<b>active-paths</b> —Consider only active paths during the unicast reverse-path check. <b>feasible-paths</b> —Consider all feasible paths during the unicast reverse-path check. <b>Default:</b> If you omit the <b>unicast-reverse-path</b> statement, only the active paths to a particular destination are considered.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks on page 105</a></li> <li>• <a href="#">Enabling Unicast Reverse-Path Forwarding Check for VPNs on page 104</a></li> </ul>

## vpn-apply-export

---

<b>Syntax</b>	vpn-apply-export;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor</i> ], [edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>neighbor</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply both the VRF export and BGP group or neighbor export policies (VRF first, then BGP) before routes from the <b>vrf</b> or <b>l2vpn</b> routing tables are advertised to other PE routers.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Applying Both the VRF Export and the BGP Export Policies on page 30</a></li></ul>

## vrf-export

<b>Syntax</b>	<code>vrf-export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols  vpls mesh-group <i>mesh-group-name</i>]  [edit routing-instances <i>routing-instance-name</i>]  [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]  [edit switch-options]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.</p>
<b>Description</b>	<p>Specify how routes are exported from the local PE router's VRF table (<i>routing-instance-name</i>.inet.0) to the remote PE router. If the value <b>vrf</b> is specified for the <b>instance-type</b> statement included in the routing instance configuration, this statement is required.</p> <p>You can configure multiple export policies on the PE router or PE switch.</p>
<b>Default</b>	If the instance-type is <b>vrf</b> , <b>vrf-export</b> is a required statement. The default action is to reject.
<b>Options</b>	<b><i>policy-names</i></b> —Names for the export policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Implementing EVPN-VXLAN for Data Centers</a></li> <li>• <a href="#">instance-type on page 1017</a></li> <li>• <a href="#">Configuring an Export Policy for the PE Router's VRF Table on page 29</a></li> </ul>

## vrf-import

---

<b>Syntax</b>	<code>vrf-import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>    vpls mesh-group <i>mesh-group-name</i>]</code> <code>[edit routing-instances <i>routing-instance-name</i>]</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]</code> <code>[edit switch-options]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.
<b>Description</b>	<p>Specify how routes are imported into the virtual routing and forwarding (VRF) table (<i>routing-instance-name</i>.inet.0) of the local provider edge (PE) router or switch from the remote PE router. If the value <b>vrf</b> is specified for the <b>instance-type</b> statement included in the routing instance configuration, this statement is required.</p> <p>You can configure multiple import policies on the PE router or switch.</p>
<b>Default</b>	If the instance type is <b>vrf</b> , <b>vrf-import</b> is a required statement. The default action is to accept.
<b>Options</b>	<i>policy-names</i> —Names for the import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Implementing EVPN-VXLAN for Data Centers</a></li><li>• <a href="#">instance-type on page 1017</a></li><li>• <a href="#">Configuring an Import Policy for the PE Router's VRF Table on page 27</a></li></ul>

---

## vrf-mtu-check

---

<b>Syntax</b>	vrf-mtu-check;
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	On M Series routers (except the M120 and M320 router), T Series routers, and on EX Series 8200 switches, configure path maximum transmission unit (MTU) checks on the outgoing interface for unicast traffic routed on a virtual private network (VPN) routing and forwarding (VRF) instance.
<b>Default</b>	Disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Path MTU Checks for VPN Routing Instances on page 16</a></li><li>• <i>Configuring the Junos OS to Enable MTU Path Check for a Routing Instance on M Series Routers</i></li></ul>

## vrf-target

<b>Syntax</b>	<pre>vrf-target {     community;     auto     import community-name;     export community-name; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn vni-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>],</p> <p>[edit switch-options]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches. <b>auto option</b> was also added at this time.</p>
<b>Description</b>	<p>Specify a virtual routing and forwarding (VRF) target community. If you configure the <b>community</b> option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. The purpose of the <b>vrf-target</b> statement is to simplify the configuration by allowing you to configure most statements at the [edit routing-instances] hierarchy level. In effect, this statement configures a single policy for import and a single policy for export to replace the per-VRF policies for every community.</p> <p>You can still create more complex policies by explicitly configuring VRF import and export policies using the <b>import</b> and <b>export</b> options.</p>
<b>Options</b>	<p><b>community</b>—Community name.</p> <p><b>auto</b>—Automatically derives the route target (RT) for QFX5100 switches.</p> <p><b>import community-name</b>—Allowed communities accepted from neighbors.</p> <p><b>export community-name</b>—Allowed communities sent to neighbors.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring a VRF Target on page 31](#)
  - *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*





## CHAPTER 50

# Configuration Statements (Layer 3 VPNs)


- [advertise-from-main-vpn-tables](#) on page 1040
- [advertise-mode \(MPLS\)](#) on page 1041
- [all-regions](#) on page 1042
- [arp-prefix-limit \(Host Fast Reroute\)](#) on page 1043
- [as-path-compare](#) on page 1044
- [chained-composite-next-hop](#) on page 1045
- [classifiers](#) on page 1047
- [create-new-ucast-tunnel](#) on page 1048
- [domain-id](#) on page 1049
- [domain-vpn-tag](#) on page 1050
- [dynamic-tunnels](#) on page 1051
- [egress-protection \(BGP\)](#) on page 1052
- [egress-protection \(MPLS\)](#) on page 1053
- [egress-protection \(Routing Instances\)](#) on page 1054
- [export-target](#) on page 1054
- [extended-space](#) on page 1055
- [family \(VRF Advertisement\)](#) on page 1056
- [forwarding-table](#) on page 1057
- [forwarding-context \(Protocols BGP\)](#) on page 1058
- [global-arp-prefix-limit \(Host Fast Reroute\)](#) on page 1059
- [global-supplementary-blackout-timer \(Host Fast Reroute\)](#) on page 1061
- [group \(Routing Instances\)](#) on page 1063
- [group-address \(Routing Instances VPN\)](#) on page 1065
- [group-range \(MBGP MVPN Tunnel\)](#) on page 1066
- [group-rp-mapping](#) on page 1067
- [host-fast-reroute](#) on page 1068
- [import-target](#) on page 1069
- [independent-domain](#) on page 1070

- [inet-mvpn \(BGP\) on page 1071](#)
- [inet-mvpn \(VRF Advertisement\) on page 1072](#)
- [inet6 on page 1072](#)
- [inet6-mvpn \(BGP\) on page 1073](#)
- [inet6-mvpn \(VRF Advertisement\) on page 1074](#)
- [inet6-vpn on page 1075](#)
- [ingress \(Chained Composite Next Hop\) on page 1076](#)
- [ingress-replication on page 1078](#)
- [interface \(Host Fast Reroute\) on page 1079](#)
- [ip-tunnel-rpf-check on page 1080](#)
- [interface \(Virtual Tunnel in Routing Instances\) on page 1081](#)
- [inter-region on page 1082](#)
- [inter-region-segmented on page 1083](#)
- [inter-region-template on page 1085](#)
- [l3vpn on page 1087](#)
- [l3vpn \(transit\) on page 1088](#)
- [label on page 1089](#)
- [label-switched-path \(Protocols MVPN\) on page 1090](#)
- [label-switched-path-template \(Multicast\) on page 1091](#)
- [labeled-bgp on page 1092](#)
- [labeled-unicast on page 1093](#)
- [link-protection \(Host Fast Reroute\) on page 1094](#)
- [localized-fib on page 1095](#)
- [maximum-paths on page 1096](#)
- [maximum-prefixes on page 1098](#)
- [metric \(Protocols OSPF Sham Link\) on page 1099](#)
- [mpls-internet-multicast on page 1100](#)
- [multicast \(Virtual Tunnel in Routing Instances\) on page 1100](#)
- [multihop on page 1101](#)
- [multipath \(Routing Options\) on page 1103](#)
- [mvpn on page 1104](#)
- [mvpn-mode on page 1106](#)
- [no-vrf-advertise on page 1107](#)
- [no-vrf-propagate-ttl on page 1108](#)
- [per-group-label on page 1109](#)
- [pim-asm on page 1109](#)
- [pim-ssm \(Selective Tunnel\) on page 1110](#)

- [primary \(Virtual Tunnel in Routing Instances\)](#) on page 1111
- [protect core](#) on page 1112
- [protection \(Protocols BGP\)](#) on page 1113
- [provider-tunnel](#) on page 1114
- [redundancy-group \(Chassis - MX Series\)](#) on page 1118
- [redundancy-group \(Interfaces\)](#) on page 1119
- [redundant-logical-tunnel](#) on page 1120
- [redundant-virtual-tunnel](#) on page 1121
- [region](#) on page 1122
- [register-limit](#) on page 1123
- [route-target \(Protocols MVPN\)](#) on page 1124
- [routing-instances \(CoS\)](#) on page 1125
- [rpt-spt](#) on page 1126
- [rsvp-te \(Routing Instances Provider Tunnel Selective\)](#) on page 1127
- [rsvp-te \(Protocols MVPN\)](#) on page 1128
- [selective](#) on page 1129
- [sglimit](#) on page 1131
- [sham-link](#) on page 1132
- [sham-link-remote](#) on page 1133
- [source \(Routing Instances Provider Tunnel Selective\)](#) on page 1134
- [source-class-usage](#) on page 1135
- [spt-only](#) on page 1136
- [static-lsp](#) on page 1137
- [supplementary-blackout-timer \(Host Fast Reroute\)](#) on page 1139
- [target \(Routing Instances MVPN\)](#) on page 1140
- [template\(Protocols MVPN\)](#) on page 1141
- [threshold-rate](#) on page 1143
- [traceoptions \(Protocols MVPN\)](#) on page 1144
- [traffic-statistics \(Protocols BGP\)](#) on page 1146
- [tunnel-limit \(Routing Instances Provider Tunnel Selective\)](#) on page 1147
- [unicast \(Route Target Community\)](#) on page 1148
- [unicast \(Virtual Tunnel in Routing Instances\)](#) on page 1149
- [vpn-localization](#) on page 1150
- [vpn-unequal-cost](#) on page 1151
- [vrf-advertise-selective](#) on page 1152
- [vrf-propagate-ttl](#) on page 1153
- [vrf-table-label](#) on page 1154

- [wildcard-group-inet on page 1156](#)
- [wildcard-group-inet6 on page 1157](#)
- [wildcard-source \(Selective Provider Tunnels\) on page 1158](#)

## advertise-from-main-vpn-tables

<b>Syntax</b>	advertise-from-main-vpn-tables;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp],
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3.
<b>Description</b>	<p>Advertise VPN routes from the main VPN tables in the master routing instance (for example, bgp.l3vpn.0, bgp.mvpn.0) instead of advertising VPN routes from the tables in the VPN routing instances (for example, <i>instance-name</i>.inet.0, <i>instance-name</i>.mvpn.0). Enable nonstop active routing (NSR) support for BGP multicast VPN (MVPN).</p> <p>When this statement is enabled, before advertising a route for a VPN prefix, the path selection algorithm is run on all routes (local and received) that have the same route distinguisher (RD).</p>
	<div>  <p><b>NOTE:</b> Adding or removing this statement causes all BGP sessions that have VPN address families to be removed and then added again. On the other hand, having this statement in the configuration prevents BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.</p> </div>
<b>Default</b>	If you do not include this statement, VPN routes are advertised from the tables in the VPN routing instances.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Junos OS Routing Tables</a></li> <li>• <a href="#">Types of VPNs on page 4</a></li> </ul>

## advertise-mode (MPLS)

<b>Syntax</b>	<code>advertise-mode (stub-alias   stub-proxy);</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls egress-protection context-identifier <i>context-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> egress-protection context-identifier <i>context-id</i>],</p> <p>[edit protocols mpls egress-protection context-identifier <i>context-id</i>],</p> <p>[edit protocols mpls label-switched-path <i>lsp-name</i> egress-protection context-identifier <i>context-id</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	<p>Configure the method for the interior gateway protocol (IGP) to advertise egress protection availability.</p> <p>Egress protection availability is advertised in the IGP. Label protocols along with CSPF use this information to do egress protection.</p>
<b>Options</b>	<p><b>stub-alias</b>—Context identifier has an alias.</p> <p>In the alias method, the LSP end-point address has an explicit backup egress node where the backup node can be learned or configured on the penultimate hop node (PHN) of a protected LSP. With this model, the PHN of a protected LSP sets up the bypass LSP tunnel to back up the egress node by avoiding the primary egress node. This model requires a Junos OS upgrade in core nodes, but is flexible enough to support all traffic engineering constraints.</p> <p><b>stub-proxy</b>—Context-identifier has a stub proxy node.</p> <p>A stub node is one that only appears at the end of an AS path, which means it does not provide transit service. In this mode, known as the virtual or proxy mode, the LSP end-point address is represented as a node with bidirectional links, with the LSP's primary egress node and backup egress node. With this representation, the penultimate hop of the LSP primary egress point can behave like a PLR in setting up a bypass tunnel to back up the egress by avoiding the primary egress node. This model has the advantage that you do not need to upgrade Junos OS on core nodes and thereby helps operators to deploy this technology.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Egress Protection for Layer 3 VPN Edge Protection Overview on page 827</a></li> <li>• <a href="#">Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP on page 843</a></li> </ul>

## all-regions

<b>Syntax</b>	<pre> all-regions {   incoming;   ingress-replication {     create-new-ucast-tunnel;     label-switched-path {       label-switched-path-template (Multicast) {         (default-template   lsp-template-name);       }     }   }   ldp-p2mp;   rsvp-te {     label-switched-path-template (Multicast) {       (default-template   lsp-template-name);     }     static-lsp static-lsp;   } } </pre>
<b>Hierarchy Level</b>	[edit protocols mvpn inter-region-template template <i>template-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Specify the tunnel support for all regions that are not defined in the more specific regions. The <b>all-regions</b> statement supports different tunnel types such as incoming, ingress replication, ldp-p2mp, or rsvp-te. The incoming tunnel indicates that the tunnel type across the area boundary router (ABR) is not changed.
<b>Options</b>	<p><b>incoming</b>— Specify the tunnel type to be the same across the ABR.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li> <li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li> <li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li> <li>• <a href="#">inter-region on page 1082</a></li> <li>• <a href="#">inter-region-segmented on page 1083</a></li> <li>• <a href="#">inter-region-template on page 1085</a></li> <li>• <a href="#">region on page 1122</a></li> <li>• <a href="#">template on page 1141</a></li> </ul>

## arp-prefix-limit (Host Fast Reroute)

<b>Syntax</b>	<code>arp-prefix-limit <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Override the ARP prefix limit for the specified host fast-reroute (HFRR) profile.</p> <p>When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.</p> <p>There are two configuration statements (<a href="#">global-arp-prefix-limit</a> and <a href="#">arp-prefix-limit</a>) that set the ARP prefix limit, one at the global <code>[edit routing-options host-fast-reroute]</code> hierarchy level and the other at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, respectively. The global <a href="#">global-arp-prefix-limit</a> statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The <a href="#">arp-prefix-limit</a> statement overrides the <a href="#">global-arp-prefix-limit</a> for that HFRR profile for that protected interface.</p> <p>Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.</p> <p>After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.</p> <p>There are global and per-HFRR CLI statements (<a href="#">global-supplementary-blackout-timer</a> and <a href="#">supplementary-blackout-timer</a>) to configure the supplementary blackout timer. The global value is at the <code>[edit routing-options host-fast-reroute]</code> hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, and overrides the global value for that HFRR profile only.</p> <p>When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.</p> <p>If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the <code>restart routing</code> command.</p>

<b>Default</b>	If you omit the <b>arp-prefix-limit</b> statement, the <b>global-arp-prefix-limit</b> takes effect for all HFRR profiles on the device. If you omit both of these statements, there is no ARP prefix limit for host fast reroute.
<b>Options</b>	<b>number</b> —Maximum number of ARP HFRR routes allowed. <b>Range:</b> 1 through 10,000 HFRR routes
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link Protection with Host Fast Reroute on page 911</a></li></ul>

---

## as-path-compare

---

<b>Syntax</b>	as-path-compare;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multipath], [edit routing-instances <i>routing-instance-name</i> routing-options multipath]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Specify to have the algorithm that is used to determine the active path compare the AS numbers in the AS path. In a VPN scenario with multiple BGP paths, the algorithm selects as the active path the route whose AS numbers match. By default, the algorithm evaluates only the length and not the contents of the AS path.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Algorithm That Determines the Active Route to Evaluate AS Numbers in AS Paths for VPN Routes on page 253</a></li></ul>



## chained-composite-next-hop

**Syntax** chained-composite-next-hop {  
     ingress;  
     transit;  
 }

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options [forwarding-table](#)],  
 [edit routing-options [forwarding-table](#)]



**NOTE:** The [edit logical-systems] hierarchy level is not supported on the QFX10000 switches.

**Release Information** Statement introduced in Junos OS Release 12.1.  
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.  
 Statement introduced in Junos OS Release 15.1 for QFX10000 Series switches.

**Description** Allows you to configure the chained composite next hops for devices handling ingress or transit traffic in the network.

Chained composite next hops help to facilitate the handling of large volumes of transit traffic in the core of large networks by allowing the router to process much larger volumes of routes. A chained composite next hop allows the router to direct sets of routes sharing the same destination to a common forwarding next hop, rather than having each route also include the destination. In the event that a network destination is changed, rather than having to update all of the routes sharing that destination with the new information, just the shared forwarding next hop is updated with the new information. The chained composite next hops continue to point to this forwarding next hop which now contains the new destination.

On platforms containing only MPCs, such as PTX Series Packet Transport Routers, the MX80 router, the MX2020 router, and the QFX10000 switches, chained composite next hops are enabled by default. On MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs and on T4000 Core Routers containing MPC and FPCs, chained composite next hops are disabled by default and need to be explicitly configured.



**NOTE:**

- Starting with Junos OS Release 13.3, for chained composite next hop feature to take effect for directly connected PE devices, the chassis must be configured to use the `enhanced-ip` option (in the case of MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs) or the `enhanced-mode` option (in the case of T4000 Core Routers containing MPC and FPCs) in the network service mode, in addition to the `l3vpn` configuration.

For more information about configuring chassis network services, see the *Junos OS Administration Library*.

- On MX Series routers, removing the `chained-composite-next-hop` statement from a PE device configuration causes all IBGP sessions to be torn down and triggers the BGP session to flap as well. A similar change on a router configured as a route reflector does not have any effect, however.

The following is a sample system log message that is generated to record such an event:

```
Nov  6 15:16:21.670 host PE1: rpd[6947]: bgp_peer_mgmt_clear:5995:
NOTIFICATION sent to 10.0.100.2 (External AS 100): code 6 (Cease)
subcode 4 (Administratively Reset), Reason: Management session cleared
BGP neighbor
```



**NOTE:** Starting in Junos OS Release 14.1, the `transit l3vpn` statement is enabled by default on PTX Series Packet Transport Routers only.

The remaining statements are explained separately. See [CLI Explorer](#).

**Default** This statement is disabled by default.

**Options** `ingress`—Enable or disable composite chained next hop for ingress traffic.  
`transit`—Enable or disable composite chained next hop for transit traffic.

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**

- [Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928](#)
- [Chained Composite Next Hops for Transit Devices for VPNs](#)
- [Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs on page 86](#)
- [ingress on page 1076](#)
- [transit \(Chained Composite Next Hops\)](#)

## classifiers

---

<b>Syntax</b>	<pre>classifiers {     exp (classifier-name   default); }</pre>
<b>Hierarchy Level</b>	[edit class-of-service routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For routing instances with VRF table labels enabled, apply a custom MPLS EXP classifier to the routing instance. You can apply the default MPLS EXP classifier or one that is previously defined.
<b>Default</b>	If you do not include this statement, the default MPLS EXP classifier is applied to the routing instance.
<b>Options</b>	<i>classifier-name</i> —Name of the behavior aggregate MPLS EXP classifier.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs on page 268</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

## create-new-ucast-tunnel

---

<b>Syntax</b>	create-new-ucast-tunnel;
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> ingress-replication]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	One of two modes for building unicast tunnels when ingress replication is configured for the provider tunnel. When this statement is configured, each time a new destination is added to the multicast distribution tree, a new unicast tunnel to the destination is created in the ingress replication tunnel. The new tunnel is deleted if the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i></li><li>• <a href="#">mpls-internet-multicast on page 1100</a></li><li>• <a href="#">ingress-replication on page 1078</a></li></ul>

## domain-id

---


<b>Syntax</b>	<code>domain-id <i>domain-id</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify a domain ID for a route. The domain ID identifies the OSPF domain from which the route originated.
<b>Options</b>	<p><b><i>domain-id</i></b>—You can specify either an IP address or an IP address and a local identifier using the following format: <b><i>ip-address:local-identifier</i></b>. If you do not specify a local identifier with the IP address, the identifier is assumed to have a value of 0.</p> <p><b>Default:</b> If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an OSPF Domain ID on page 174</a></li> </ul>

## domain-vpn-tag

---

<b>Syntax</b>	<code>domain-vpn-tag <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) routing device.
<b>Options</b>	<i>number</i> —VPN tag.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring an OSPF Domain ID on page 174</a></li></ul>

## dynamic-tunnels

<b>Syntax</b>	<pre>dynamic-tunnels <i>tunnel-name</i> {   destination-networks <i>prefix</i>;   gre;   rsvp-te <i>entry-name</i> {     destination-networks <i>network-prefix</i>;     label-switched-path-template (Multicast) {       default-template;       template-name;     }   }   source-address <i>address</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
<b>Description</b>	Configure a dynamic tunnel between two PE routers.
<div>  <b>NOTE:</b> ACX Series routers do not support the <code>gre</code> statement. </div>	
<b>Options</b>	<p><b><i>tunnel-name</i></b>—Name of the dynamic tunnel.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks</i></li> </ul>

## egress-protection (BGP)

---

<b>Syntax</b>	<pre>egress-protection {     context-identifier <i>context-id-ip-address</i>;     keep-import <i>policy-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> family inet-vpn unicast], [edit protocols bgp group <i>group-name</i> family inet6-vpn unicast], [edit protocols bgp group <i>group-name</i> family iso-vpn unicast]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4R3. Statement introduced for labeled unicast in Junos OS Release 14.1.
<b>Description</b>	Enable egress protection for the configured BGP Network Layer Reachability Information (NLRI).
<b>Options</b>	<p><b>context-identifier <i>context-id-ip-address</i></b>—(Optional) The IPv4 address of the context identifier.</p> <p><b>keep-import <i>policy-name</i></b>—The import policy that contains routes with the configured route target community. BGP keeps all these routes in the VPN routing table. The protector PE router scans the policy and builds the remote instance with the configured community name from the policy.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Egress Protection for Layer 3 VPN Services on page 834</a></li><li>• <a href="#">Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP on page 843</a></li><li>• <a href="#">Example: Configuring Egress Protection for BGP Labeled Unicast on page 813</a></li></ul>



## egress-protection (MPLS)

<b>Syntax</b>	<pre>egress-protection {   context-identifier <i>context-id</i> {     primary   protector;     metric <i>igp-metric-value</i>;     advertise-mode (stub-alias   stub-proxy);   } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Options <b>primary</b> , <b>protector</b> , and <b>metric</b> introduced in Junos OS Release 11.4R3. Option <b>advertise-mode</b> introduced in Junos OS Release 13.3.
<b>Description</b>	Enables an Edge Protection Virtual Circuit (EPVC) for the MPLS protocol.
<b>Options</b>	<p><b>context-identifier <i>context-id-ip-address</i></b>—(Optional) The context identifier IPv4 address.</p> <p><b>metric <i>igp-metric-value</i></b>—(Optional) The IGP metric value ranging from 2 through 16777215.</p> <p><b>(primary   protector)</b>—On the primary PE router, configure as type <b>primary</b>. On the protector PE router, configure as type <b>protector</b>.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Egress Protection for Layer 3 VPN Services on page 834</a></li> <li>• <a href="#">Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP on page 843</a></li> </ul>

## egress-protection (Routing Instances)

---



<b>Syntax</b>	<pre>egress-protection {     context-identifier <i>context-id-ip-address</i>; }</pre>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> ].
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4R3.
<b>Description</b>	Enable egress instance protection and provides customer edge (CE) VRF-level context-id granularity for each virtual routing and forwarding (VRF) instance.
<b>Options</b>	<b>context-identifier <i>context-id-ip-address</i></b> —(Optional) The IPv4 address for the context identifier of the protected PE router.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Egress Protection for Layer 3 VPN Services on page 834</a></li></ul>

## export-target

---

<b>Syntax</b>	<pre>export-target {     target <i>target-community</i>;     unicast; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN network layer reachability information (NLRI).
<b>Options</b>	<b>target <i>target-community</i></b> —Specify the export target community. <b>unicast</b> —Use the same target community as specified for unicast.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## extended-space

<b>Syntax</b>	extended-space;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table chained-composite-next-hop ingress <a href="#">l3vpn</a> ], [edit routing-options forwarding-table chained-composite-next-hop ingress <a href="#">l3vpn</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Accept more than one million Layer 3 VPN BGP updates with unique inner VPN labels. The neighboring PE routers are typically from other vendors and configured to assign a unique inner label to each Layer 3 VPN BGP route.
	<div>  <p><b>NOTE:</b> This statement is supported on Juniper Networks routers containing only MPCs and MICs and with chassis network services configured with the <b>enhanced-ip</b> setting.</p> </div>
	<div>  <p><b>WARNING:</b> You cannot configure the <b>extended-space</b> statement on a MX system with a MS-DPC unless you have also configured FIB localization with the MS-DPC set as a FIB-remote. This is because the MS-DPC cannot install nexthop in <b>extended-space</b> so the use of a default nexthop pointing to a FIB-local FPC is necessary to ensure proper data forwarding.</p> </div>
<b>Default</b>	This statement is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928</a></li> </ul>

## family (VRF Advertisement)

---

<b>Syntax</b>	<pre>family {     inet-mvpn;     inet6-mvpn; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">vrf-advertise-selective</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">vrf-advertise-selective</a> ],
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	<p>Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.</p> <p>The options are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring PIM-SSM GRE Selective Provider Tunnels on page 597</a></li><li>• <a href="#">inet-mvpn on page 1072</a></li><li>• <a href="#">inet6-mvpn on page 1074</a></li></ul>

## forwarding-table

<b>Syntax</b>	<pre>forwarding-table {   chained-composite-next-hop;   ecmp-fast-reroute,   export [ <i>policy-name</i> ];   (indirect-next-hop   no-indirect-next-hop);   (indirect-next-hop-change-acknowledgements      no-indirect-next-hop-change-acknowledgements;)   krt-nexthop-ack-timeout <i>interval</i>;   unicast-reverse-path (active-paths   feasible-paths); }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
<b>Release Information</b>	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 14.1X53-D30 for the QFX Series.
<b>Description</b>	Configure information about the routing device's forwarding table.  The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Load Balancing BGP Traffic</i></li> </ul>

## forwarding-context (Protocols BGP)

---

Syntax	forwarding-context;
Hierarchy Level	[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ],
Release Information	Statement introduced in Junos OS Release 16.1 for the M Series, MX Series, and T Series.
Description	<p>Routing instance used for data forwarding and transport session. Use this statement in conjunction with <b>mpls-forwarding</b>, to protect against label spoofing across AS boundary routers in the context of Inter-AS VPN Option B AS boundary routers. If <b>forwarding-context</b> is not set for a VPN BGP peer, both the routing instance and forwarding context are provided by the master routing instance.</p> <p>Setting a neighbor interface as a forwarding context can be useful, for example, when configuring a common AS boundary router so that it only accepts MPLS packets from a peer AS boundary router whose labels were explicitly advertised to the common AS boundary router. The action creates a special routing instance (MFI) that is used to form a transport session with peer and to perform forwarding functions for traffic from the peer. Spoof checking is performed for peers with different MFI forwarding contexts. Because peers with the same forwarding context share the same <b>MFI.mpls.0</b> table, there is no need for spoof checking on these peers. Thus, in <b>foo.mpls.0</b>, BGP installs only VPN labels that have been advertised to the VPN AS boundary router peers using the MFI <b>foo</b> as the forwarding context. This ensures that only MPLS packets from the VPN AS boundary router that also have the advertised label are accepted.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <i>Example: Preventing BGP Session Resets</i></li><li>• <i>Junos OS Administration Library</i></li></ul>

## global-arp-prefix-limit (Host Fast Reroute)

<b>Syntax</b>	<code>global-arp-prefix-limit <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-options host-fast-reroute]</code> , <code>[edit routing-options host-fast-reroute]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Set the ARP prefix limit for all host fast-reroute (HFRR) profiles on the routing device.</p> <p>When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.</p> <p>There are two configuration statements (<code>global-arp-prefix-limit</code> and <code>arp-prefix-limit</code>) that set the ARP prefix limit, one at the global <code>[edit routing-options host-fast-reroute]</code> hierarchy level and the other at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, respectively. The global <code>global-arp-prefix-limit</code> statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The <code>arp-prefix-limit</code> statement overrides the <code>global-arp-prefix-limit</code> for that HFRR profile for that protected interface.</p> <p>Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.</p> <p>After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.</p> <p>There are global and per-HFRR CLI statements (<code>global-supplementary-blackout-timer</code> and <code>supplementary-blackout-timer</code>) to configure the supplementary blackout timer. The global value is at the <code>[edit routing-options host-fast-reroute]</code> hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, and overrides the global value for that HFRR profile only.</p> <p>When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.</p> <p>If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the <code>restart routing</code> command.</p>

**Default** If you omit the **-arp-prefix-limit** statement, the **global-arp-prefix-limit** takes effect for all HFRR profiles on the device. If you omit both of these statements, there is no ARP prefix limit for host fast reroute.

**Options** *number*—Maximum number of ARP HFRR routes allowed.  
**Range:** 1 through 10,000 HFRR routes

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Link Protection with Host Fast Reroute on page 911](#)



## global-supplementary-blackout-timer (Host Fast Reroute)

<b>Syntax</b>	<code>global-supplementary-blackout-timer <i>minutes</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-options host-fast-reroute],</code> <code>[edit routing-options host-fast-reroute]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Set the blackout timer for all host fast-reroute (HFRR) profiles on the routing device.</p> <p>When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.</p> <p>There are two configuration statements (<code>global-arp-prefix-limit</code> and <code>arp-prefix-limit</code>) that set the ARP prefix limit, one at the global <code>[edit routing-options host-fast-reroute]</code> hierarchy level and the other at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, respectively. The global <code>global-arp-prefix-limit</code> statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The <code>arp-prefix-limit</code> statement overrides the <code>global-arp-prefix-limit</code> for that HFRR profile for that protected interface.</p> <p>Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.</p> <p>After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.</p> <p>There are global and per-HFRR CLI statements (<code>global-supplementary-blackout-timer</code> and <code>supplementary-blackout-timer</code>) to configure the supplementary blackout timer. The global value is at the <code>[edit routing-options host-fast-reroute]</code> hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, and overrides the global value for that HFRR profile only.</p> <p>When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.</p> <p>If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the <code>restart routing</code> command.</p>
<b>Default</b>	If you omit the <code>supplementary-blackout-timer</code> statement, the <code>global-supplementary-blackout-timer</code> takes effect for all HFRR profiles on the device. If

you omit both of these statements, the blackout timeout value is set to the ARP cache timeout value, which by default is 20 minutes and is configurable with the **aging-timer** statement at the **[edit system arp]** hierarchy level.

**Options**    **minutes**—Number of minutes, in addition to the ARP cache timeout value, before all HFRR profiles on the routing device are reactivated after the ARP prefix limit is exceeded.

**Range:** 1 through 15 minutes

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                         routing-control—To add this statement to the configuration.

**Related Documentation**    • [Example: Configuring Link Protection with Host Fast Reroute on page 911](#)

## group (Routing Instances)

```
Syntax  group address {
        source source-address {
            inter-region-segmented {
                fan-out fan-out value;
                threshold rate-value;
            }
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
        wildcard-source {
            inter-region-segmented {
                fan-out fan-out value;
            }
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
        }
    }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective],  
[edit routing-instances *routing-instance-name* provider-tunnel selective]

**Release Information** Statement introduced in Junos OS Release 8.5.  
The **inter-region-segmented** statement added in Junos OS Release 15.1.

**Description** Specify the IP address for the multicast group configured for point-to-multipoint label-switched paths (LSPs) and PIM-SSM GRE selective provider tunnels.

**Options** **address**—Specify the IP address for the multicast group. This address must be a valid multicast group address.


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**

- [Configuring the Multicast Group Address for an MBGP MVPN on page 468](#)
- [Configuring PIM-SSM GRE Selective Provider Tunnels on page 597](#)

## group-address (Routing Instances VPN)

<b>Syntax</b>	<code>group-address address;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel pim-asm],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel pim-asm family inet],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel pim-asm family inet6],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel pim-ssm],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel pim-ssm family inet],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel pim-ssm family inet6],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel pim-asm],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel pim-asm family inet],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel pim-asm family inet6],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel pim-ssm],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel pim-ssm family inet],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel pim-ssm family inet6]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Starting with Junos OS Release 11.4, to provide consistency with draft-rosen 7 and next-generation BGP-based multicast VPNs, configure the provider tunnels for draft-rosen 6 anysource multicast VPNs at the [edit routing-instances <i>routing-instance-name</i> provider-tunnel] hierarchy level. The <b>mdt</b>, <b>vpn-tunnel-source</b>, and <b>vpn-group-address</b> statements are deprecated at the [edit routing-instances <i>routing-instance-name</i> protocols pim] hierarchy level. Use <b>group-address</b> in place of <b>vpn-group-address</b>.</p>
<b>Description</b>	Specify a group address on which to encapsulate multicast traffic from a virtual private network (VPN) instance.
<div>  <p><b>NOTE:</b> IPv6 provider tunnels are not currently supported for draft-rosen MVPNs. They are supported for MBGP MVPNs.</p> </div>	
<b>Options</b>	<p><b>address</b>—For IPv4, IP address whose high-order bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. For IPv6, IP address whose high-order bits are FF00 (FF00::/8).</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Example: Configuring Any-Source Multicast for Draft-Rosen VPNs</i></li> </ul>

- [Configuring Multicast Layer 3 VPNs on page 611](#)
- [Multicast Protocols Feature Guide](#)


---

## group-range (MBGP MVPN Tunnel)

---

<b>Syntax</b>	<code>group-range <i>multicast-prefix</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> source <i>source-address</i> pim-ssm],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source pim-ssm],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm],</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> source <i>source-address</i> pim-ssm],</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source pim-ssm],</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm],</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.
<b>Options</b>	<b><i>multicast-prefix</i></b> —Multicast group address range to be used to create MBGP MVPN source-specific multicast selective PMSI tunnels. <b>Range:</b> Any valid, nonreserved IPv4 multicast address range <b>Default:</b> None
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.

## group-rp-mapping

<b>Syntax</b>	<pre>group-rp-mapping {   family (inet   inet6) {     log-interval <i>seconds</i>;     maximum <i>limit</i>;     threshold <i>value</i>;   }   log-interval <i>seconds</i>;   maximum <i>limit</i>;   threshold <i>value</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],  [edit protocols pim rp],  [edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	Configure a limit for the number of incoming group-to-RP mappings.
<div>  <p><b>NOTE:</b> The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet   inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum group-to-RP mapping limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>	
<b>Options</b>	<p><b>family (inet   inet6)</b>—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured group-to-RP mapping limit.</p> <p><b>Default:</b> Both IPv4 and IPv6 messages are counted towards the configured group-to-RP limit.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM State Limits on page 545</a></li> </ul>

## host-fast-reroute

---



<b>Syntax</b>	<pre>host-fast-reroute {     global-arp-prefix-limit <i>number</i>;     global-supplementary-blackout-timer <i>minutes</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Configure global settings for all host fast reroute (HFRR) profiles configured on a routing device.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link Protection with Host Fast Reroute on page 911</a></li></ul>



## import-target

<b>Syntax</b>	<pre> import-target {   target {     target-value;     receiver target-value;     sender target-value;   }   unicast {     receiver;     sender;   } } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## independent-domain

<b>Syntax</b>	<code>independent-domain &lt;no-attrset&gt;;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options autonomous-system <i>autonomous-system</i> ], [edit routing-instances <i>routing-instance-name</i> routing-options autonomous-system <i>autonomous-system</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>no-attrset</b> option introduced in Junos OS Release 10.4.
<b>Description</b>	<p>Configure an independent AS domain.</p> <p>The independent domain uses transitive path attribute 128 (attribute set) messages to tunnel the independent domain's BGP attributes through the internal BGP (IBGP) core.</p> <p>This improves the transparency of Layer 3 VPN services for customer networks by preventing the IBGP routes that originate within an autonomous system (AS) in the customer network from being sent to a service provider's AS. Similarly, IBGP routes that originate within an AS in the service provider's network are prevented from being sent to a customer AS.</p>
	<p> <b>NOTE:</b> In Junos OS Release 10.3 and later, if BGP receives attribute 128 and you have not configured an independent domain in any routing instance, BGP treats the received attribute 128 as an unknown attribute.</p>
	<p> <b>NOTE:</b> The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.</p>
<b>Options</b>	<b>no-attrset</b> —(Optional) Disables attribute set messages on the independent AS domain.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Tunneling Layer 3 VPN IPv6 Islands over an IPv4 Core Using IBGP and Independent Domains on page 231</a></li> <li>• <a href="#">Configuring Layer 3 VPNs to Carry IBGP Traffic on page 241</a></li> <li>• <i>autonomous-system</i></li> </ul>

## inet-mvpn (BGP)

<b>Syntax</b>	<pre>inet-mvpn {   signaling {     accepted-prefix-limit {       maximum <i>number</i>;       teardown <i>percentage</i> {         idle-timeout (forever   <i>minutes</i>);       }     }   }   damping;   loops <i>number</i>;   prefix-limit {     maximum <i>number</i>;     teardown <i>percentage</i> {       idle-timeout (forever   <i>minutes</i>);     }   } }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family],  [edit protocols bgp family],  [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family],  [edit protocols bgp group <i>group-name</i> family]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	Enable the <b>inet-mvpn</b> address family in BGP.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>

## inet-mvpn (VRF Advertisement)

---

<b>Syntax</b>	inet-mvpn;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">vrf-advertise-selective family</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">vrf-advertise-selective family</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	Enable IPv4 MVPN routes to be advertised from the VRF instance.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Limiting Routes to Be Advertised by an MVPN VRF Instance on page 590</a></li></ul>

## inet6

---

<b>Syntax</b>	inet6;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table chained-composite-next-hop ingress <a href="#">labeled-bgp</a> ], [edit routing-options forwarding-table chained-composite-next-hop ingress <a href="#">labeled-bgp</a> ]
<b>Release Information</b>	This statement was introduced in Junos OS Release 12.3.
<b>Description</b>	Allows you to configure chained composite next hops for IPv6 labeled-unicast routes.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928</a></li></ul>

## inet6-mvpn (BGP)

```
Syntax  inet6-mvpn {
        signaling {
            accepted-prefix-limit {
                maximum number;
                teardown percentage {
                    idle-timeout (forever | minutes);
                }
            }
        }
        loops number
        prefix-limit {
            maximum number;
            teardown percentage {
                idle-timeout (forever | minutes);
            }
        }
    }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp family],  
[edit protocols bgp family],  
[edit logical-systems *logical-system-name* protocols bgp group *group-name* family],  
[edit protocols bgp group *group-name* family]

**Release Information** Statement introduced in Junos OS Release 10.0.

**Description** Enable the **inet6-mvpn** address family in BGP.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.


**Related Documentation** • *BGP Configuration Overview*

## inet6-mvpn (VRF Advertisement)

---

<b>Syntax</b>	inet6-mvpn;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">vrf-advertise-selective family</a> ], [edit routing-instances <i>routing-instance-name</i> <a href="#">vrf-advertise-selective family</a> ],
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	Enable IPv6 MVPN routes to be advertised from the VRF instance.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## inet6-vpn

<b>Syntax</b>	inet6-vpn (any   multicast   unicast) { aggregate-label; prefix-limit <i>maximum</i> ; rib-group <i>rib-group-name</i> ; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols bgp family], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family], [edit protocols bgp family], [edit protocols bgp group <i>group-name</i> family]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Enable IP version 6 (IPv6) on the provider edge (PE) router for the Layer 3 VPN.
<div>  <p><b>NOTE:</b> The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.</p> </div>	
<b>Options</b>	<p><b>any</b>—Configure the family type to be both multicast and unicast.</p> <p><b>multicast</b>—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.</p> <p><b>prefix-limit <i>maximum</i></b>—Maximum prefix limit.  <b>Range:</b> 1 through 4,294,967,295  <b>Default:</b> 1</p> <p><b>rib-group <i>rib-group-name</i></b>—The name of the routing table group.</p> <p><b>unicast</b>—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Layer 3 VPNs to Carry IPv6 Traffic on page 227</a></li> <li>• <i>Junos OS Routing Protocols Library</i></li> </ul>

## ingress (Chained Composite Next Hop)

**Syntax**    ingress {  
               (evpn | no-evpn);  
               (fec129-vpws | no-fec129-vpws);  
               (l2ckt | no-l2ckt);  
               (l2vpn | no-l2vpn);  
               l3vpn;  
               labeled-bgp;  
           }

**Hierarchy Level**    [edit logical-systems *logical-system-name* routing-options forwarding-table  
                           **chained-composite-next-hop**],  
                           [edit routing-options forwarding-table **chained-composite-next-hop**]

**Release Information**    This statement was introduced in Junos OS Release 12.1.  
                               **labeled-bgp** option was introduced in Junos OS Release 12.3.  
                               **l2ckt**, **l2vpn**, **no-l2ckt**, and **no-l2vpn** options were introduced in Junos OS Release 13.3.  
                               **evpn**, **fec129-vpws**, **no-evpn**, and **no-fec129-vpws** options were introduced in Junos OS Release 14.1.

**Description**    Allows you to configure chained composite next hops for devices handling transit traffic in the network.

The other statements are explained separately.



**NOTE:** The **extended-space** statement is not supported in ACX Series routers.



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

**Default**    This statement is disabled by default.

**Options**    **evpn | no-evpn**—Enable or disable composite chained next hop for ingress EVPN label-switched paths (LSPs).

**fec129-vpws | no-fec129-vpws**—Enable or disable composite chained next hop for ingress FEC 129 virtual private wire service (VPWS) LSPs.

**l2ckt | no-l2ckt**—Enable or disable composite chained next hop for ingress Layer 2 circuit LSPs.

**l2vpn | no-l2vpn**—Enable or disable composite chained next hop for ingress Layer 2 virtual private network (VPN) LSPs.



The remaining statements are explained separately. See [CLI Explorer](#).

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928</a></li><li>• <a href="#">chained-composite-next-hop on page 1045</a></li></ul>
------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## ingress-replication

<b>Syntax</b>	<pre>ingress-replication {   create-new-ucast-tunnel;   label-switched-path {     label-switched-path-template {       (template-name   default-template);     }   } }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel],</p> <p>[edit protocols mvpn inter-region-template template <i>template-name</i> all-regions],</p> <p>[edit protocols mvpn inter-region-template template <i>template-name</i> region <i>region-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	<p>A provider tunnel type used for passing multicast traffic between routers through the MPLS cloud, or between PE routers when using MVPN. The ingress replication provider tunnel uses MPLS point-to-point LSPs to create the multicast distribution tree.</p> <p>Optionally, you can specify a label-switched path template. If you configure <b>ingress-replication label-switched-path</b> and do not include <b>label-switched-path-template</b>, ingress replication works with existing LDP or RSVP tunnels. If you include <b>label-switched-path-template</b>, the tunnels must be RSVP.</p>
<b>Options</b>	<p><b>existing-unicast-tunnel</b>—An existing tunnel to the destination is used for ingress replication. If an existing tunnel is not available, the destination is not added. Default mode if no option is specified.</p> <p><b>create-new-ucast-tunnel</b>—When specified, a new unicast tunnel to the destination is created and used for ingress replication. The unicast tunnel is deleted later if the destination is no longer included in the multicast distribution tree.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i></li> <li>• <a href="#">create-new-ucast-tunnel on page 1048</a></li> <li>• <a href="#">mpls-internet-multicast on page 1100</a></li> </ul>

## interface (Host Fast Reroute)

---

<b>Syntax</b>	<pre>interface <i>interface-name</i> {   arp-prefix-limit <i>number</i>;   link-protection;   supplementary-blackout-timer <i>minutes</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Configure host fast reroute settings, including link protection for the interface and an ARP prefix limit.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Link Protection with Host Fast Reroute on page 911</a></li> </ul>

## ip-tunnel-rpf-check

---

<b>Syntax</b>	<pre>ip-tunnel-rpf-check {     mode (<i>strict</i>   <i>loose</i>);     fail-filter <i>filter-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> routing-options forwarding-table]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.1 for MX Series Routers with MICs.
<b>Description</b>	<p>Configure the system to enable anti-spoofing protection for next-hop-based dynamic tunnels, where reverse path forwarding checks are placed to ensure that the tunnel traffic is received from a legitimate source through designated IP tunnel, where the source is reachable on the same tunnel on which the packet was received.</p> <p>When a packet comes from a nondesignated source, the reverse path forwarding check fails in the strict mode, and passes in the loose mode. When a packet comes from a nonexistent source, the reverse path forwarding check fails.</p> <p>By default, the reverse path forwarding check is in strict mode, where the packets are not forwarded if the source of the packet is from a nondesignated tunnel.</p>
<b>Options</b>	<p><b>mode (<i>strict</i>   <i>loose</i>)</b>—(Optional) Specify the mode of the reverse path forwarding check to enable anti-spoofing protection for next-hop-based dynamic tunnels.</p> <p>In the strict mode (default), the reverse path forwarding check fails when the packet is received from a nondesignated tunnel source. The check passes only for packets from designated tunnels.</p> <p>In the loose mode, the reverse path forwarding check passes even if the packet is received from a nondesignated tunnel source.</p> <p>When the packet is from a nonexistent tunnel source, the reverse path forwarding check fails in both the strict and loose modes.</p> <p><b>Default:</b> If you omit the mode statement, the default behavior is strict mode.</p> <p><b>fail-filter <i>filter-name</i></b>—(Optional) Attach a filter to the Layer 3 VPN to log packets that failed the reverse path forwarding check.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview on page 753</a></li><li>• <a href="#">Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels on page 756</a></li></ul>

## interface (Virtual Tunnel in Routing Instances)

<b>Syntax</b>	<pre>interface vt-<i>fpc/pic/port.unit-number</i> {     multicast;     primary;     unicast; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4.
<b>Description</b>	<p>In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure a virtual tunnel (VT) interface.</p> <p>VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic).</p> <p>In an MBGP MVPN extranet, if there is more than one VRF routing instance on a PE router that has receivers interested in receiving multicast traffic from the same source, VT interfaces must be configured on all instances.</p> <p>Starting in Junos OS Release 12.3, you can configure multiple VT interfaces in each routing instance. This provides redundancy. A VT interface can be used in only one routing instance.</p>
<b>Options</b>	<p><i>vt-fpc/pic/port.unit-number</i>—Name of the VT interface.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 600</a></li> <li>• <a href="#">Example: Configuring MBGP MVPN Extranets</a></li> </ul>

## inter-region

---

<b>Syntax</b>	<pre>inter-region [     template <i>template-name</i>;     no-inter-region-segmentation; ]</pre>
<b>Hierarchy Level</b>	[edit routing-instances <i>instance-name</i> provider-tunnel]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Specify template to be used in inter-region segmentation. The tunnels types used will be based on the template mentioned. If a virtual routing and forwarding (VRF) in the area boundary router (ABR) does not want to participate in the inter-region segmentation then <b>no-inter-region-segmentation</b> can be specified.
<b>Options</b>	<p><b>no-inter-region-segmentation</b>— Specify to avoid the template mapping an ABR that is not involved in the inter-region segmentation.</p> <p><b>template <i>template-name</i></b>— Specify the template to be used in the inter-region segmentation.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li><li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li><li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li><li>• <a href="#">inter-region-segmented on page 1083</a></li><li>• <a href="#">inter-region-template on page 1085</a></li><li>• <a href="#">region on page 1122</a></li><li>• <a href="#">template on page 1141</a></li></ul>

## inter-region-segmented

<b>Syntax</b>	<pre>inter-region-segmented {     fan-out <i>fan-out value</i>;     threshold <i>rate-value</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit routing-instances <i>instance-name</i> provider-tunnel selective group <i>mcast</i> IP address source <i>source IP-address</i>], [edit routing-instances <i>instance-name</i> provider-tunnel selective group <i>mcast</i> IP address wildcard-source], [edit routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source], [edit routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	<p>Configure segmentation of inter-region on an ingress PE router or transit area boundary routers (ABRs) by specifying only a fan-out factor or threshold rate or both. The inter-region segmentation allows the PE router to signal segmented S-PMSI A-D routes. If the segmentation is based on the threshold, then the MVPN starts migrating the traffic flow to a segmented S-PMSI when the threshold configured for the selective tunnel reaches the threshold rate. The threshold value can be configured for S-PMSI with LDP-P2MP, IR, and RSVP-TE tunnel types. If the fan-out value is configured, then the traffic flow is moved to segmented S-PMSI when the number of leaf A-D routes per S-PMSI meets the configured fan-out value. The fan-out value can be configured for wildcard S-PMSI, and S-PMSI.</p> <p>If both threshold and fan-out values are set for S-PMSI with LDP-P2MP, IR, and RSVP-TE tunnel types, then the traffic flow to the segmented inter-region happens only when the traffic rate multiplied by the number of leaf A-D routes exceeds the threshold value.</p>
<b>Options</b>	<p><b>fan-out <i>fan-out-value</i></b>— Specify the number of remote leaf A-D route values for fan-out. Once the number of leaf A-D routes per S-PMSI meets the configured fan-out value, the inter-region segmentation is triggered and the traffic flow is directed to segmented S-PMSI.</p> <p><b>Range:</b> 1 through 10000</p> <p><b>threshold <i>rate-value</i></b>— Specify the data threshold rate to trigger inter-region segmentation.</p>



**NOTE:** This option is available only at the following hierarchy levels:

```
[edit routing-instances instance-name provider-tunnel selective group
mcast IP address source source IP-address],
[edit routing-instances instance-name provider-tunnel selective group
mcast IP address wildcard-source]
```

**Range:** 0 through 1000000 Kilobits

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li><li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li><li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li><li>• <a href="#">all-regions on page 1042</a></li><li>• <a href="#">inter-region on page 1082</a></li><li>• <a href="#">inter-region-template on page 1085</a></li><li>• <a href="#">region on page 1122</a></li><li>• <a href="#">template on page 1141</a></li></ul>
------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------



## inter-region-template

```
Syntax  inter-region-template {
        template template-name {
            all-regions {
                incoming;
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template (Multicast) {
                            (default-template | lsp-template-name);
                        }
                    }
                }
            }
            ldp-p2mp;
            rsvp-te {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
            }
            static-lsp static-lsp;
            region region-name {
                incoming;
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template (Multicast) {
                            (default-template | lsp-template-name);
                        }
                    }
                }
            }
            ldp-p2mp;
            rsvp-te {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
            }
            static-lsp static-lsp;
        }
    }
}
```

**Hierarchy Level** [edit protocols mvpn]

**Release Information** Statement introduced in Junos OS Release 15.1.

**Description** Define MVPN inter-region tunnel mapping template. A template is configured on a transit area boundary router (ABR) to define the tunnel type to be used for a given region. There are four types of tunnels that can be used:

- Incoming — Indicates that the tunnel type used across the ABR remains unchanged.
- Ingress-replication — A tunnel type used for passing multicast traffic between routers through the MPLS cloud, or between PE routers when using MVPN. The ingress

replication provider tunnel uses MPLS point-to-point LSPs to create the multicast distribution tree.

- `Ldp-p2mp` — Specify a point-to-multipoint provider tunnel with LDP signaling for an MBGP MVPN.
- `Rsvp-te` — Configure RSVP tunnels to multicast traffic flooding using point-to-multipoint LSPs.

The remaining statements are explained separately. See [CLI Explorer](#).

<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li><li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li><li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li><li>• <a href="#">all-regions on page 1042</a></li><li>• <a href="#">inter-region-segmented on page 1083</a></li><li>• <a href="#">region on page 1122</a></li><li>• <a href="#">template on page 1141</a></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## l3vpn

<b>Syntax</b>	<code>l3vpn {     extended-space; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table chained-composite-next-hop <a href="#">ingress</a> ], [edit routing-options forwarding-table chained-composite-next-hop <a href="#">ingress</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	(M120 routers, M320 routers with Enhanced III FPCs, MX Series routers, and T Series routers only) Accept larger numbers of Layer 3 VPN BGP updates with unique inner VPN labels (up to one million). The neighboring PE routers are typically from other vendors and are configured to assign a unique inner label to each Layer 3 VPN BGP route.



### NOTE:

- On MX Series routers, this statement is not supported when the router has both DPCs and MPCs installed. You must remove one or the other type of card to successfully use this statement.
- In Junos OS Release 11.4 and earlier, the `l3vpn` statement was titled as `l3vpn-composite-nexthop`, and the statement was available at the [edit logical-systems *logical-system-name* routing-options] and [edit routing-options] hierarchy levels.
- On ACX Series routers, the `extended-space` statement is not supported.



**NOTE:** The [edit logical-systems] hierarchy level is not applicable in ACX Series routers.

The remaining statement is explained separately.

<b>Default</b>	This statement is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928</a></li> </ul>

- [extended-space on page 1055](#)

## **l3vpn (transit)**

---

<b>Syntax</b>	transit l3vpn;
<b>Hierarchy Level</b>	[edit logical-systems logical-system-name routing-options forwarding-table chained-composite-next-hop transit] [edit routing-options forwarding-table chained-composite-next-hop transit]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Configure l3vpn settings on the transit router.



**NOTE:** You can configure transit l3vpn only after configuring the enhanced-ip operation mode. See *network-services* to configure the enhanced-ip mode.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">chained-composite-next-hop on page 1045</a>

## label

---

<b>Syntax</b>	<pre>label {     allocation <i>label-allocation-policy</i>;     substitution <i>label-substitution-policy</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options]</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	Specify label allocation and substitution policies on a per-route basis.
<b>Options</b>	<p><b>allocation</b> <i>label-allocation-policy</i>—Specify a policy to generate labels on a per-route basis.</p> <p><b>substitution</b> <i>label-substitution-policy</i>—Specify a policy to substitute labels on a per-route basis. The label substitution policy is used to determine whether or not a label should be substituted on an AS border router. The results of the policy operation are either accept (label substitution is performed) or reject (label substitution is not performed).</p> <p><b>Default:</b> accept</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Label Allocation and Substitution Policy for VPNs on page 166</a></li> </ul>

## label-switched-path (Protocols MVPN)

---

<b>Syntax</b>	<pre>label-switched-path {     label-switched-path-template (Multicast){         (default-template   lsp-template-name);     } }</pre>
<b>Hierarchy Level</b>	<pre>[edit protocols mvpn inter-region-template template <i>template-name</i> region <i>region-name</i>   ingress-replication] [edit protocols mvpn inter-region-template template <i>template-name</i> all-regions   ingress-replication]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Specify point-to-point LSP unicast tunnel at the ABRs. Optionally, you can specify a label-switched path template. If you configure <b>ingress-replication label-switched-path</b> and do not include <b>label-switched-path-template</b> , ingress replication works with existing LDP or RSVP tunnels. If you include label-switched-path-template, then the RSVP tunnels should be enabled across the ABR regions.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li><li>• <a href="#">all-regions on page 1042</a></li><li>• <a href="#">inter-region on page 1082</a></li><li>• <a href="#">inter-region-template on page 1085</a></li><li>• <a href="#">inter-region-segmented on page 1083</a></li><li>• <a href="#">region on page 1122</a></li><li>• <a href="#">template on page 1141</a></li><li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li><li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li></ul>

## label-switched-path-template (Multicast)

<b>Syntax</b>	<pre>label-switched-path-template {   (default-template   <i>lsp-template-name</i>); }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel ingress-replication label-switched-path], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>   rsvp-te <i>entry-name</i>], [edit protocols mvpn inter-region-segmented template <i>template-name</i> region <i>region-name</i>   ingress-replication label-switched-path], [edit protocols mvpn inter-region-segmented template <i>template-name</i> region <i>region-name</i>   rsvpe-te], [edit protocols mvpn inter-region-template template <i>template-name</i> all-regions   ingress-replication label-switched-path], [edit protocols mvpn inter-region-template template <i>template-name</i> all-regions rsvp-te], [edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication   label-switched-path], [edit routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source   <i>source-address</i> rsvp-te], [edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>] [edit routing-instances <i>instance-name</i> provider-tunnel]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 18.2. under the heirarchy level [edit routing-instances <i>instance-name</i> provider-tunnel]</p>
<b>Description</b>	<p>Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the <b>label-switched-path-template</b> statement, so you must configure either the default-template using the <b>default-template</b> option, or you must specify the name of your preconfigured LSP template.</p>
<b>Options</b>	<p><b>default-template</b>—Specify that the default LSP template be used for the dynamically generated LSPs.</p> <p><b><i>lsp-template-name</i></b>—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</a></li><li>• <a href="#">Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 465</a></li><li>• <a href="#">Configuring Dynamic Point-to-Multipoint Flooding LSPs</a></li><li>• <a href="#">Configuring RSVP Automatic Mesh</a></li></ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

---


## labeled-bgp

---

Syntax	labeled-bgp { <a href="#">inet6</a> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table chained-composite-next-hop <a href="#">ingress</a> ], [edit routing-options forwarding-table chained-composite-next-hop <a href="#">ingress</a> ]
Release Information	This statement was introduced in Junos OS Release 12.3.
Description	Allows you to configure chained composite next hops for IPv6 labeled-unicast routes.  The other statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• <a href="#">Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs on page 928</a>



## labeled-unicast

<b>Syntax</b>	<pre>labeled-unicast {   aggregate-label {     community <i>community-name</i>;   }   explicit-null {     connected-only;   }   per-group-label;   resolve-vpn;   traffic-statistics {     file <i>filename</i> &lt;world-readable   no-world-readable&gt;;     interval <i>seconds</i>;   } }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet],</p> <p>[edit protocols bgp family inet],</p> <p>[edit protocols bgp group <i>group-name</i> family inet]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.</p> <p>Statement introduced in Junos OS Release 17.1 for QFX10000 switches.</p>
<b>Description</b>	<p>Advertise labeled routes from the inet.0 VPN, and place labeled routes into the inet.0 VPN. When the <b>labeled-unicast</b> statement is used, the local router automatically performs a next hop to self on all routes advertised into EBGp from IBGP and into IBGP from EBGp.</p>
<b>Options</b>	<p><b>resolve-vpn</b>—(Optional) Store labeled routes in the inet.3 or inet6.3 routing table to resolve routes for a provider edge (PE) router located in a different autonomous system (AS). For a PE router to install a route in the virtual routing and forwarding (VRF) table, the next hop must resolve to a route stored in the inet.3 or inet6.3 routing table. This option is also used to configure inter-AS VPLS with MAC operations.</p>
<div>  <p><b>NOTE:</b> VPLS is not supported by QFX Series switches.</p> </div> <p>The other statements are explained separately.</p>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service on page 402</a></li> </ul>

- [Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service on page 407](#)
- [Configuring Inter-AS VPLS with MAC Processing at the ASBR](#)

---

## link-protection (Host Fast Reroute)

---

<b>Syntax</b>	link-protection;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> routing-options interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	Enable link protection on a specified interface in a routing instance.  Configuring link protection on an interface ensures that traffic is rerouted to the destination device through alternate loop-free paths that traverse the protected interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Example: Configuring Link Protection with Host Fast Reroute on page 911</a>

## localized-fib

---

<b>Syntax</b>	localized-fib { fpc-slot <i>fpc-slot-number</i> ; }
<b>Hierarchy Level</b>	[edit routing-instances <i>instance-name</i> routing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2.
<b>Description</b>	Localize the VRF routing instance routes to a specific FPC of CE-facing physical interfaces, or specify an FPC slot number if the CE-facing interfaces are logical interfaces like AE or RSQL or IRB.
<b>Options</b>	<b>fpc-slot <i>fpc-slot-number</i></b> — Specify the FPC slot number to localize the VRF routing instance for logical interfaces like AE or RSQL or IRB.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding VRF Localization in Layer 3 VPNs on page 145</a></li> <li>• <a href="#">Example: Improving Scalability Using VRF Localization for Layer 3 VPNs on page 147</a></li> <li>• <a href="#">vpn-localization on page 1150</a></li> </ul>

## maximum-paths

**Syntax** `maximum-paths path-limit <log-interval seconds> <log-only | threshold value>;`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced in Junos OS Release 8.0.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a limit for the number of routes installed in a routing table based upon the route path.



**NOTE:** The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

**Options** `log-interval seconds`—(Optional) Minimum time interval (in seconds) between log messages.

**Range:** 5 through 86,400

`log-only`—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

`path-limit`—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.

**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ )

**Default:** No default

`threshold value`—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the `path-limit` value that starts triggering the warnings.

**Range:** 1 through 100



**NOTE:** When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *path-limit* value, then additional routes are rejected.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs on page 260</a></li></ul>
------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------

## maximum-prefixes

**Syntax** `maximum-prefixes prefix-limit <log-interval seconds> <log-only | threshold percentage>;`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced in Junos OS Release 8.0.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a limit for the number of routes installed in a routing table based upon the route prefix.

Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.



**NOTE:** The `maximum-prefixes` statement is similar to the `maximum-paths` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

**Options** `log-interval seconds`—(Optional) Minimum time interval (in seconds) between log messages.

**Range:** 5 through 86,400

`log-only`—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

`prefix-limit`—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.

**Range:** 1 through 4,294,967,295

**Default:** No default

**threshold value**—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the *prefix-limit* value that starts triggering the warnings.

**Range:** 1 through 100



**NOTE:** When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *prefix-limit* value, then additional routes are rejected.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs on page 260](#)

## metric (Protocols OSPF Sham Link)

**Syntax** *metric number;*

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id* sham-link-remote *address*],  
[edit routing-instances *routing-instance-name* protocols ospf area *area-id* sham-link-remote *address*]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Specify the cost of using the OSPF sham link.

**Default** *number*—1

**Options** *number*—1 through 65,535

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Giving VPN Backbones Priority with OSPFv2 Sham Links for Layer 3 VPNs on page 185](#)

## mpls-internet-multicast

<b>Syntax</b>	<code>mpls-internet-multicast;</code>
<b>Hierarchy Level</b>	<code>[edit routing-instances <i>routing-instance-name</i> instance-type]</code> <code>[edit protocols pim]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	<p>A nonforwarding routing instance type that supports Internet multicast over an MPLS network for the default master instance. No interfaces can be configured for it. Only one <b>mpls-internet-multicast</b> instance can be configured for each logical system.</p> <p>The <b>mpls-internet-multicast</b> configuration statement is also explicitly required under PIM in the master instance.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i></li> <li>• <a href="#">ingress-replication on page 1078</a></li> </ul>

## multicast (Virtual Tunnel in Routing Instances)

<b>Syntax</b>	<code>multicast;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interface vt-<i>fpc/pic/port.unit-number</i>]</code> , <code>[edit routing-instances <i>routing-instance-name</i> interface vt-<i>fpc/pic/port.unit-number</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4.
<b>Description</b>	In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used for multicast traffic only.
<b>Default</b>	If you omit this statement, the VT interface can be used for both multicast and unicast traffic.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 600</i></li> <li>• <i>Example: Configuring MBGP MVPN Extranets</i></li> </ul>



## multihop

<b>Syntax</b>	<pre>multihop {   no-nexthop-change;   ttl <i>ttl-value</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>   neighbor <i>address</i>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure an EBGp multihop session.</p> <p>For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.</p> <p>An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.</p> <p>If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.</p>



**NOTE:** You cannot configure the `accept-remote-nexthop` statement at the same time.

**Default** If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.


The remaining statements are explained separately. See [CLI Explorer](#).

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring EBGp Multihop Sessions*
- [Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs on page 195](#)
- *accept-remote-nexthop*
- *no-nexthop-change*
- *tll*

## multipath (Routing Options)

<b>Syntax</b>	<pre> multipath {   vpn-unequal-cost equal-external-internal;   as-path-compare; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>equal-external-internal</b> option added in Junos OS Release 8.4.</p> <p><b>as-path-compare</b> option introduced in Junos OS Release 10.1</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p>
<b>Description</b>	<p>Enable protocol-independent load balancing for Layer 3 VPNs. This allows the forwarding next hops for both the active route and alternative paths to be used for load balancing. For IPv6, you must configure the <b>multipath</b> statement at both the [edit routing-instances <i>routing-instance-name</i> routing-options] hierarchy level and the [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i>] hierarchy level.</p> <p>Protocol-independent load balancing is applied to VPN routes that are equal up to their router identifiers (<b>router-id</b>) with regard to route selection.</p> <p>The options are explained separately.</p>
	<div>  <p><b>NOTE:</b> ACX Series routers do not support <b>vpn-unequal-cost</b> statement.</p> </div>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Protocol-Independent Load Balancing in Layer 3 VPNs on page 787</a></li> </ul>

## mvpn

```
Syntax  mvpn {
    inter-region-template{
        template template-name {
            all-regions {
                incoming;
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template (Multicast) {
                            (default-template | lsp-template-name);
                        }
                    }
                }
            }
            ldp-p2mp;
            rsvp-te {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
            }
            static-lsp static-lsp;
            region region-name{
                incoming;
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path {
                        label-switched-path-template (Multicast){
                            (default-template | lsp-template-name);
                        }
                    }
                }
            }
            ldp-p2mp;
            rsvp-te {
                label-switched-path-template (Multicast) {
                    (default-template | lsp-template-name);
                }
            }
            static-lsp static-lsp;
        }
    }
}

mvpn-mode (rpt-spt | spt-only);
receiver-site;
sender-site;
route-target {
    export-target {
        target target-community;
        unicast;
    }
    import-target {
        target {
            target-value;
            receiver target-value;
            sender target-value;
        }
    }
}
```

```

    }
    unicast {
        receiver;
        sender;
    }
}
}
}

```


<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for the <b>traceoptions</b> statement at the [edit protocols mvpn] hierarchy level introduced in Junos OS Release 13.3. Support for the <b>inter-region-template</b> statement at the [edit protocols mvpn] hierarchy level introduced in Junos OS Release 15.1.
<b>Description</b>	Enable next-generation multicast VPNs in a routing instance.
<b>Options</b>	<p><b>receiver-site</b>—Allow sites with multicast receivers.</p> <p><b>sender-site</b>—Allow sites with multicast senders.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Routing Instances for an MBGP MVPN on page 537</a></li> </ul>

## mvpn-mode

---

<b>Syntax</b>	mvpn-mode ( <a href="#">rpt-spt</a>   <a href="#">spt-only</a> );
<b>Hierarchy Level</b>	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols mvpn], [edit routing-instances <i>instance-name</i> protocols mvpn]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	Configure the mode for customer PIM (C-PIM) join messages. Mixing MVPN modes within the same VPN is not supported. For example, you cannot have spt-only mode on a source PE and rpt-spt mode on the receiver PE.
<b>Default</b>	spt-only
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 538</a></li><li>• <a href="#">Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 540</a></li></ul>

## no-vrf-advertise

<b>Syntax</b>	no-vrf-advertise;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	<p>Prevent advertising VPN routes from a VRF routing instance to remote PE routers.</p> <p>Within a hub-and-spoke configuration, you can configure a PE router to not advertise VPN routes from the primary (hub) routing instance. Instead, these routes are advertised from the secondary (downstream) routing instance. You can do this without configuring routing table groups, by using the <b>no-vrf-advertise</b> statement.</p>
<div>  <p><b>NOTE:</b> This statement does not prevent the exportation of VPN routes to other VRF routing instances on the same router by configuring the [edit routing-options auto-export] statement.</p> </div>	
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a VRF Target on page 31</a></li> <li>• <a href="#">Configuring Hub-and-Spoke VPN Topologies: One Interface on page 653</a></li> <li>• <a href="#">Limiting Routes to Be Advertised by an MVPN VRF Instance on page 590</a></li> <li>• <a href="#">vrf-advertise-selective on page 1152</a></li> </ul>

## no-vrf-propagate-ttl

---

<b>Syntax</b>	no-vrf-propagate-ttl;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit protocols routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Disable normal time-to-live (TTL) decrementing in a VRF routing instance. You configure this statement once per routing instance, and it affects only RSVP-signaled or LDP-signaled LSPs in the routing instance. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.
<b>Default</b>	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing on page 989</a></li><li>• <a href="#">Disabling Normal TTL Decrementing</a></li></ul>



## per-group-label

---

<b>Syntax</b>	<code>per-group-label;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> family inet labeled-unicast]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.
<b>Description</b>	Account for traffic from each customer separately by advertising separate labels for the same prefix to the peer routers in the BGP groups.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics on page 314</a></li> <li>• <a href="#">Understanding Interprovider and Carrier-of-Carriers VPNs on page 312</a></li> </ul>

## pim-asm

---

<b>Syntax</b>	<code>pim-asm {   group-address (Routing Instances) <i>address</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel], [edit routing-instances <i>routing-instance-name</i> provider-tunnel]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3.
<b>Description</b>	Specify a Protocol Independent Multicast (PIM) sparse mode provider tunnel for an MBGP MVPN or for a draft-rosen MVPN.  The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## pim-ssm (Selective Tunnel)

---

<b>Syntax</b>	<pre>pim-ssm {     group-range <i>multicast-prefix</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel selective group <i>group-address</i> source <i>source-address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel selective group <i>group-address</i> wildcard-source], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel selective wildcard-group-inet wildcard-source], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   provider-tunnel selective wildcard-group-inet6 wildcard-source], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i>   source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i>   wildcard-source], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet   wildcard-source], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6   wildcard-source]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## primary (Virtual Tunnel in Routing Instances)

<b>Syntax</b>	<code>primary;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3.
<b>Description</b>	<p>In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used as the primary interface for multicast traffic.</p> <p>Junos OS supports up to eight VT interfaces configured for multicast in a routing instance to provide redundancy for MBGP (next-generation) MVPNs. This support is for RSVP point-to-multipoint provider tunnels as well as multicast Label Distribution Protocol (MLDP) provider tunnels. This feature works for extranets as well.</p> <p>This statement allows you to configure one of the VT interfaces to be the primary interface, which is always used if it is operational. If a VT interface is configured as the primary, it becomes the nexthop that is used for traffic coming in from the core on the label-switched path (LSP) into the routing instance. When a VT interface is configured to be primary and the VT interface is used for both unicast and multicast traffic, only the multicast traffic is affected.</p> <p>If no VT interface is configured to be the primary or if the primary VT interface is unusable, one of the usable configured VT interfaces is chosen to be the nexthop that is used for traffic coming in from the core on the LSP into the routing instance. If the VT interface in use goes down for any reason, another usable configured VT interface in the routing instance is chosen. When the VT interface in use changes, all multicast routes in the instance also switch their reverse-path forwarding (RPF) interface to the new VT interface to allow the traffic to be received.</p> <p>To realize the full benefit of redundancy, we recommend that when you configure multiple VT interfaces, at least one of the VT interfaces be on a different Tunnel PIC from the other VT interfaces. However, Junos OS does not enforce this.</p>
<b>Default</b>	If you omit this statement, Junos OS chooses a VT interface to be the active interface for multicast traffic.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 600</a></li> <li>• <a href="#">Example: Configuring MBGP MVPN Extranets</a></li> </ul>

## protect core

---

<b>Syntax</b>	protect core;
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> routing-options]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 13.2.</p> <p>Support in Layer 3 VPN routers is extended to BGP PIC with multiple routes in the global tables such as inet and inet6 unicast, and inet and inet6 labeled unicast in Junos OS Release 15.1.</p> <p>Support for Resource Reservation Protocol (RSVP) introduced in Junos OS Release 16.1.</p>
<b>Description</b>	<p>In an MPLS VPN Core network, enable BGP Prefix-Independent Convergence (PIC) Edge to install a VPN route in the forwarding table for use as an alternate path when a provider edge (PE) router fails or connectivity to a PE router is lost. This alternate link is used until global convergence through the interior gateway protocol (IGP) occurs. Both OSPF with LDP or RSVP, and IS-IS with LDP or RSVP are supported.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BGP PIC Edge for MPLS Layer 3 VPNs on page 791</a></li><li>• <a href="#">Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 878</a></li><li>• <a href="#">Introduction to Configuring Layer 3 VPNs on page 130</a></li><li>• <i>LDP Operation</i></li><li>• <i>Understanding Multiprotocol BGP</i></li></ul>

---

## protection (Protocols BGP)

---

<b>Syntax</b>	protection;
<b>Hierarchy Level</b>	[edit routing-instances <i>instance-name</i> protocols bgp family inet unicast], [edit routing-instances <i>instance-name</i> protocols bgp family inet6 unicast], [edit routing-instances <i>instance-name</i> protocols bgp family inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp family inet6 labeled-unicast]
<b>Description</b>	Configure the backup path to protect the active provider edge path in a Layer 3 VPN or a BGP labeled unicast path.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Provider Edge Link Protection in Layer 3 VPNs on page 878</a></li><li>• <a href="#">Example: Configuring Provider Edge Link Protection for BGP Labeled Unicast Paths on page 895</a></li></ul>

## provider-tunnel

```

Syntax  provider-tunnel {
        family {
            inet {
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                }
            }
            ldp-p2mp;
            mdt {
                data-mdt-reuse;
                group-range multicast-prefix;
                threshold {
                    group group-address {
                        source source-address {
                            rate threshold-rate;
                        }
                    }
                }
            }
            tunnel-limit limit;
        }
        pim-asm {
            group-address (Routing Instances) address;
        }
        pim-ssm {
            group-address (Routing Instances) address;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
    }
    inet6 {
        ingress-replication {
            create-new-ucast-tunnel;
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
        ldp-p2mp;
        mdt {
            data-mdt-reuse;
            group-range multicast-prefix;
            threshold {
                group group-address {
                    source source-address {
                        rate threshold-rate;
                    }
                }
            }
        }
    }

```

```

    }
    tunnel-limit limit;
  }
}
pim-asm {
  group-address (Routing Instances) address;
}
pim-ssm {
  group-address (Routing Instances) address;
}
rsvp-te {
  label-switched-path-template {
    (default-template | lsp-template-name);
  }
  static-lsp lsp-name;
}
ingress-replication {
  create-new-ucast-tunnel;
  label-switched-path-template {
    (default-template | lsp-template-name);
  }
}
ldp-p2mp;
pim-asm {
  group-address (Routing Instances) address;
}
pim-ssm {
  group-address (Routing Instances) address;
}
rsvp-te {
  label-switched-path-template {
    (default-template | lsp-template-name);
  }
  static-lsp lsp-name;
}
selective {
  group mcast-prefix/prefix-length {
    source ip-prefix/prefix-length {
      ldp-p2mp;
      create-new-ucast-tunnel;
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
    }
  }
  pim-ssm {
    group-range mcast-prefix;
  }
  rsvp-te {
    label-switched-path-template {
      (default-template | lsp-template-name);
    }
    static-lsp point-to-multipoint-lsp-name;
  }
  threshold-rate kpbs;
}
wildcard-source {

```

```

    pim-ssm {
        group-range multicast-prefix;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kpbs;
}
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
wildcard-group-inet6 {
    wildcard-source {
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
}
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],  
[edit routing-instances *routing-instance-name*]



<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>The <b>selective</b> statement and substatements added in Junos OS Release 8.5.</p> <p>The <b>ingress-replication</b> statement and substatements added in Junos OS Release 10.4.</p> <p>In Junos OS Release 17.3R1, the <b>mdt</b> hierarchy was moved from <b>provider-tunnel</b> to the <b>provider-tunnel family inet</b> and <b>provider-tunnel family inet6</b> hierarchies as part of an upgrade to add IPv6 support for default MDT in Rosen 7, and data MDT for Rosen 6 and Rosen 7. The <b>provider-tunnel mdt</b> hierarchy is now hidden for backward compatibility with existing scripts.</p>
<b>Description</b>	<p>Configure virtual private LAN service (VPLS) flooding of unknown unicast, broadcast, and multicast traffic using point-to-multipoint LSPs. Also configure point-to-multipoint LSPs for MBGP MVPNs.</p>
<b>Options</b>	<p>The remaining statements are explained separately. .</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Flooding Unknown Traffic Using Point-to-Multipoint LSPs in VPLS</i></li><li>• <a href="#">Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 465</a></li><li>• <i>Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode</i></li></ul>

## redundancy-group (Chassis - MX Series)

---

**Syntax**

```
redundancy-group {  
  interface-type {  
    redundant-logical-tunnel {  
      device count;  
    }  
    redundant-virtual-tunnel {  
      device count;  
    }  
  }  
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 13.3.

**Description** Configure redundant logical tunnels, redundant virtual tunnels, or both on MX Series 3D Universal Edge Routers.

The remaining statements are explained separately. See [CLI Explorer](#).

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview on page 567](#)
  - [Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 569](#)
  - [Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 570](#)
  - [redundancy-group \(Interfaces\) on page 1119](#)

## redundancy-group (Interfaces)

<b>Syntax</b>	<pre> redundancy-group {   member-interface <i>interface-name</i> {     (active   backup);   } } </pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Configure member tunnels of redundant logical or virtual tunnels only on MX Series 3D Universal Edge Routers.
<b>Options</b>	<p><b>active</b>—Set the interface to the active mode.</p> <p><b>backup</b>—Set the interface to the backup mode.</p> <p>The remaining statement is explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To view this statement in the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview on page 567</a></li> <li>• <a href="#">Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 569</a></li> <li>• <a href="#">Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 570</a></li> <li>• <a href="#">Example: Configuring Redundant Logical Tunnels</a></li> <li>• <a href="#">Configuring Redundant Logical Tunnels</a></li> <li>• <a href="#">Redundant Logical Tunnels Overview</a></li> <li>• <a href="#">redundancy-group on page 1118</a></li> </ul>

## redundant-logical-tunnel

---

<b>Syntax</b>	<code>redundant-logical-tunnel {     device-count <i>count</i>; }</code>
<b>Hierarchy Level</b>	[edit chassis redundancy-group interface-type]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3. Support for up to 255 redundant logical tunnels added to Junos OS Release 13.3R3 and 14.1R2.
<b>Description</b>	Configure redundant logical tunnels on MX Series 3D Universal Edge Routers.
<b>Options</b>	<b><i>count</i></b> —Specify the number of the redundant logical tunnels. For Junos OS Release 13.3R1, 13.3R2, and 14.1R1 the valid range is from 1 through 16. For Junos OS Release 13.3R3 and 14.1R2 and later releases the valid range is from 1 through 255.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To view this statement in the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Redundant Logical Tunnels</i></li><li>• <i>Configuring Redundant Logical Tunnels</i></li><li>• <i>Redundant Logical Tunnels Overview</i></li><li>• <a href="#">redundancy-group (Interfaces) on page 1119</a></li><li>• <a href="#">redundancy-group (Chassis - MX Series) on page 1118</a></li></ul>

## redundant-virtual-tunnel


---

<b>Syntax</b>	redundant-virtual-tunnel { device-count <i>count</i> ; }
<b>Hierarchy Level</b>	[edit chassis redundancy-group interface-type]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.2.
<b>Description</b>	Configure redundant virtual tunnels on MX Series 3D Universal Edge Routers.
<b>Options</b>	<b>device-count</b> <i>count</i> —Specify the number of redundant virtual tunnels. <b>Range:</b> 1 through 16
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Redundant Virtual Tunnels Providing Resiliency in Delivering Multicast Traffic Overview on page 567</a></li> <li>• <a href="#">Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 569</a></li> <li>• <a href="#">Example: Configuring Redundant Virtual Tunnels to Provide Resiliency in Delivering Multicast Traffic on page 570</a></li> <li>• <a href="#">redundancy-group (Interfaces) on page 1119</a></li> <li>• <a href="#">redundancy-group (Chassis - MX Series) on page 1118</a></li> </ul>

## region

<b>Syntax</b>	<pre> region <i>region-name</i>{   incoming;   ingress-replication {     create-new-ucast-tunnel;     label-switched-path {       llabel-switched-path-template (Multicast){         (default-template   <i>lsp-template-name</i>);       }     }   }   ldp-p2mp;   rsvp-te {     label-switched-path-template (Multicast) {       (default-template   <i>lsp-template-name</i>);     }     static-lsp <i>static-lsp</i>;   } } </pre>
<b>Hierarchy Level</b>	[edit protocols mvpn inter-region-template template <i>template-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Specify the regions based on the BGP peer groups. The regions support different tunnel types such as incoming, ingress replication, ldp-p2mp, or rsvp-te. The incoming tunnel indicates that the tunnel type across the area boundary router (ABR) is not changed. A template is used to define the tunnel type for each region.
<b>Options</b>	<p><b>incoming</b>— Specify the tunnel type to be same across area boundary router (ABR).</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li> <li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li> <li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li> <li>• <a href="#">all-regions on page 1042</a></li> <li>• <a href="#">inter-region on page 1082</a></li> <li>• <a href="#">inter-region-segmented on page 1083</a></li> <li>• <a href="#">inter-region-template on page 1085</a></li> <li>• <a href="#">template on page 1141</a></li> </ul>

## register-limit

<b>Syntax</b>	<pre> register-limit {     family (inet   inet6) {         log-interval <i>seconds</i>;         maximum <i>limit</i>;         threshold <i>value</i>;     }     log-interval <i>seconds</i>;     maximum <i>limit</i>;     threshold <i>value</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	Configure a limit for the number of incoming (S,G) PIM registers.
<div>  <p><b>NOTE:</b> The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet   inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum PIM register message limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>	
<b>Options</b>	<p><b>family (inet   inet6)</b>—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured register message limit.</p> <p><b>Default:</b> Both IPv4 and IPv6 messages are counted towards the configured register message limit.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM State Limits on page 545</a></li> <li>• <code>clear pim join</code></li> <li>• <code>clear pim register</code></li> </ul>

## route-target (Protocols MVPN)

<b>Syntax</b>	<pre> route-target {   export-target {     target <i>target-community</i>;     unicast;   }   import-target {     target {       target <i>target-value</i>;       receiver <i>target-value</i>;       sender <i>target-value</i>;     }     unicast {       receiver;       sender;     }   } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvpn]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.
<b>Default</b>	The multicast VPN routing instance uses the import and export route targets configured for the Layer 3 VPN.
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring VRF Route Targets for Routing Instances for an MBGP MVPN on page 587</a></li> </ul>



## routing-instances (CoS)

<b>Syntax</b>	<pre> routing-instances <i>routing-instance-name</i> {   classifiers {     dscp (<i>classifier-name</i>   default);     dscp-ipv6 (<i>classifier-name</i>   default);     exp (<i>classifier-name</i>   default);     ieee-802.1 (<i>classifier-name</i>   default);     no-default;   }   policy-map <i>policy-map-name</i>{   } } </pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>no-default</b> and <b>policy-map</b> options added for MX Series devices only in Junos OS Release 16.1.</p>
<b>Description</b>	<p>For routing instances with VRF table labels enabled, apply a DSCP, MPLS EXP, or IEEE 802.1 classifier to the routing instance. You can apply the default classifier or one that is previously defined.</p> <p>Apply the <b>no-default</b> option to disable the application of any default classifier to the routing instance.</p> <p>You can also apply a policy map that defines the rewrite rules for a routing instance.</p>
<b>Default</b>	If you do not include this statement with the <b>classifiers</b> option, the default MPLS EXP classifier is applied to the routing instance. When no DSCP classifier is configured, the default MPLS EXP classifier is applied.
<b>Options</b>	<p><b><i>routing-instance-name</i></b>—Name of the routing instance.</p> <p><b><i>classifier-name</i></b>—Name of the behavior aggregate MPLS EXP classifier or DSCP classifier.</p> <p><b><i>policy-map-name</i></b>—Name of the policy map that defines the rewrite rules for a routing instance.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Custom Forwarding Class for Each Queue</a></li> <li>• <a href="#">Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs on page 268</a></li> </ul>

## rpt-spt

---

<b>Syntax</b>	<code>rpt-spt;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols mvpn mvpn-mode], [edit routing-instances <i>instance-name</i> protocols mvpn mvpn-mode]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	Use rendezvous-point trees for customer PIM (C-PIM) join messages, and switch to the shortest-path tree after the source is known.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## rsvp-te (Routing Instances Provider Tunnel Selective)

**Syntax**

```
rsvp-te {
    label-switched-path-template {
        (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
}
```

**Hierarchy Level**

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source],  
 [edit routing-instances *routing-instance-name* provider-tunnel],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** Configure the properties of the RSVP traffic-engineered point-to-multipoint LSP for MBGP MVPNs.

The remaining statements are explained separately. See [CLI Explorer](#).



**NOTE:** Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.

**Required Privilege Level**

routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Selective Provider Tunnels for an MBGP MVPN on page 466](#)

## rsvp-te (Protocols MVPN)

---

<b>Syntax</b>	<pre>rsvp-te{   label-switched-path-template (Multicast){     (default-template   lsp-template-name);   }   static-lsp static-lsp; }</pre>
<b>Hierarchy Level</b>	[edit protocols mvpn inter-region-template template <i>template-name</i> all-regions], [edit protocols mvpn inter-region-template template <i>template-name</i> region <i>region-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Configure to unicast, broadcast, and multicast traffic flooding using point-to-multipoint LSPs.
<b>Options</b>	<b>static-lsp <i>static-lsp</i></b> — Create a static point-to-multipoint LSP and automatically include all of the neighbors in the routing instance.  The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471</a></li><li>• <a href="#">all-regions on page 1042</a></li><li>• <a href="#">inter-region on page 1082</a></li><li>• <a href="#">inter-region-template on page 1085</a></li><li>• <a href="#">inter-region-segmented on page 1083</a></li><li>• <a href="#">region on page 1122</a></li><li>• <a href="#">template on page 1141</a></li><li>• <a href="#">Example: Configuring Segmented Inter-Area P2MP LSP on page 478</a></li><li>• <a href="#">Configuring Segmented Inter-Area P2MP LSP on page 476</a></li></ul>

## selective

```

Syntax  selective {
        group multicast-prefix/prefix-length {
            source ip-prefix/prefix-length {
                ingress-replication {
                    create-new-ucast-tunnel;
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                }
            }
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp point-to-multipoint-lsp-name;
            }
            threshold-rate kbits;
        }
        wildcard-source {
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
            }
            static-lsp point-to-multipoint-lsp-name;
        }
        threshold-rate kbits;
    }
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
        static-lsp lsp-name;
    }
    threshold-rate number;
}
}

```


```

wildcard-group-inet6 {
  wildcard-source {
    ldp-p2mp;
    pim-ssm {
      group-range multicast-prefix;
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
    threshold-rate number;
  }
}

```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel], [edit routing-instances <i>routing-instance-name</i> provider-tunnel]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. The <b>ingress-replication</b> statement and substatements added in Junos OS Release 10.4.
<b>Description</b>	Configure selective point-to-multipoint LSPs for an MBGP MVPN. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the MBGP MVPNs, helping to minimize flooding in the service provider's network.  The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Selective Provider Tunnels for an MBGP MVPN on page 466</a></li> <li>• <a href="#">Configuring PIM-SSM GRE Selective Provider Tunnels on page 597</a></li> </ul>

## sglimit

<b>Syntax</b>	<pre>sglimit {     family (inet   inet6) {         log-interval <i>seconds</i>;         maximum <i>limit</i>;         threshold <i>value</i>;     }     log-interval <i>seconds</i>;     maximum <i>limit</i>;     threshold <i>value</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim ],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	Configure a limit for the number of accepted (*G) and (S,G) PIM join states.
	<div>  <p><b>NOTE:</b> The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet   inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum PIM join state limit, you cannot configure a limit at the family level for IPv4 or IPv6 joins. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>
<b>Options</b>	<p><b>family (inet   inet6)</b>—(Optional) Specify either IPv4 or IPv6 join states to be counted towards the configured join state limit.</p> <p><b>Default:</b> Both IPv4 and IPv6 join states are counted towards the configured join state limit.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM State Limits on page 545</a></li> <li>• <code>clear pim join</code></li> </ul>

## sham-link

---

<b>Syntax</b>	<pre>sham-link {     local <i>address</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf], [edit routing-instances <i>routing-instance-name</i> protocols ospf]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Configure the local endpoint of a sham link.</p> <p>You can create an intra-area link or sham link between two provider edge (PE) routing devices so that the VPN backbone is preferred over the back-door link. A back-door link is a backup link that connects customer edge (CE) devices in case the VPN backbone is unavailable. When such a backup link is available and the CE devices are in the same OSPF area, the default behavior is to prefer this backup link over the VPN backbone. This is because the backup link is considered an intra-area link, while the VPN backbone is always considered an inter-area link. Intra-area links are always preferred over inter-area links.</p> <p>The sham link is an unnumbered point-to-point intra-area link between PE devices. When the VPN backbone has a sham intra-area link, this sham link can be preferred over the backup link if the sham link has a lower OSPF metric than the backup link.</p> <p>The sham link is advertised using Type 1 link-state advertisements (LSAs). Sham links are valid only for routing instances and OSPFv2.</p> <p>Each sham link is identified by the combination of a local endpoint address and a remote endpoint address.</p>
<b>Options</b>	<b>local <i>address</i></b> —The address for the local endpoint of the sham link.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring OSPFv2 Sham Links on page 186</a></li><li>• <a href="#">sham-link-remote on page 1133</a></li></ul>



## sham-link-remote

<b>Syntax</b>	<pre>sham-link-remote address {     demand-circuit;     ipsec-sa name;     metric metric; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for <b>ipsec-sa</b> statement added in Junos OS Release 8.3.</p>
<b>Description</b>	<p>Configure the remote endpoint of a sham link.</p> <p>You can create an intra-area link or sham link between two provider edge (PE) routing devices so that the VPN backbone is preferred over the back-door link. A back-door link is a backup link that connects customer edge (CE) devices in case the VPN backbone is unavailable. When such a backup link is available and the CE devices are in the same OSPF area, the default behavior is to prefer this backup link over the VPN backbone. This is because the backup link is considered an intra-area link, while the VPN backbone is always considered an inter-area link. Intra-area links are always preferred over inter-area links.</p> <p>The sham link is an unnumbered point-to-point intra-area link between PE devices. When the VPN backbone has a sham intra-area link, this sham link can be preferred over the backup link if the sham link has a lower OSPF metric than the backup link.</p> <p>The sham link is advertised using Type 1 link-state advertisements (LSAs). Sham links are valid only for routing instances and OSPFv2.</p> <p>Each sham link is identified by the combination of a local endpoint address and a remote endpoint address.</p>
<b>Options</b>	<p><b>address</b>—Address for the remote end point of the sham link.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring OSPFv2 Sham Links on page 186</a></li> <li>• <a href="#">sham-link on page 1132</a></li> </ul>

## source (Routing Instances Provider Tunnel Selective)

<b>Syntax</b>	<pre> source <i>source-address</i> {   ldp-p2mp;   pim-ssm {     group-range <i>multicast-prefix</i>;   }   rsvp-te {     label-switched-path-template {       (default-template   <i>lsp-template-name</i>);     }     static-lsp <i>lsp-name</i>;   }   threshold-rate <i>number</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Specify the IP address for the multicast source. This statement is a part of the point-to-multipoint LSP and PIM-SSM GRE selective provider tunnel configuration for MBGP MVPNs.
<b>Options</b>	<p><b><i>source-address</i></b>—IP address for the multicast source.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Multicast Source Address for an MBGP MVPN on page 468</a></li> <li>• <a href="#">Configuring PIM-SSM GRE Selective Provider Tunnels on page 597</a></li> </ul>

## source-class-usage

<b>Syntax</b>	<pre>source-class-usage {     direction; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit routing-instances <i>routing-instance-name</i> vrf-table-label]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for the <b>vrf-table-label</b> statement added in Junos OS Release 9.3.
<b>Description</b>	Enable packet counters on an interface that count packets that arrive from specific prefixes on the provider core router and are destined for specific prefixes on the customer edge router.
<b>Options</b>	<p><i>direction</i> can be one of the following:</p> <p><b>input</b>—Configure at least one expected ingress point.</p> <p><b>output</b>—Configure at least one expected egress point.</p> <p><b>input output</b>—On a single interface, configure at least one expected ingress point and one expect egress point.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Enabling Source Class and Destination Class Usage</i></li> <li>• <i>accounting</i></li> <li>• <i>destination-class-usage</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> <li>• <a href="#">vrf-table-label on page 1154</a></li> </ul>

## spt-only

---

<b>Syntax</b>	spt-only;
<b>Hierarchy Level</b>	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols mvpn mvpn-mode], [edit routing-instances <i>instance-name</i> protocols mvpn mvpn-mode]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	Set the MVPN mode to learn about active multicast sources using multicast VPN source-active routes. This is the default mode.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 540</a></li></ul>

## static-lsp

**Syntax** `static-lsp lsp-name;`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* rsvp-te],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source rsvp-te],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source rsvp-te],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source rsvp-te],  
 [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* rsvp-te],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source rsvp-te],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source rsvp-te],  
 [edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source rsvp-te]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** Specify the name of the static point-to-multipoint (P2MP) LSP used for a specific MBGP MVPN; static P2MP LSP cannot be shared by multiple VPNs. Use this statement to specify the static LSP for both inclusive and selective point-to-multipoint LSPs.

Use a static P2MP LSP when you know all the egress PE router endpoints (receiver nodes) and you want to avoid the setup delay incurred by dynamically created P2MP LSPs (configured with the **label-switched-path-template**). These static LSPs are signaled before the MVPN requires or uses them, consequently avoiding any signaling latency and minimizing traffic loss due to latency.

If you add new endpoints after the static P2MP LSP is established, you must update the configuration on the ingress PE router. In contrast, a dynamic P2MP LSP learns new endpoints without any configuration changes.



**BEST PRACTICE:** Multiple multicast flows can share the same static P2MP LSP; this is the preferred configuration when the set of egress PE router endpoints on the LSP are all interested in the same set of multicast flows. When the set of relevant flows is different between endpoints, we recommend that you create a new static P2MP LSP to associate endpoints with flows of interest.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <i>Point-to-Multipoint LSPs Overview</i>
	• <i>Configuring Static LSPs</i>
	• <a href="#">Configuring Selective Provider Tunnels for an MBGP MVPN on page 466</a>
	• <i>Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems</i>

## supplementary-blackout-timer (Host Fast Reroute)

<b>Syntax</b>	<code>supplementary-blackout-timer <i>minutes</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Override the global supplementary blackout timer for the specified host fast-reroute (HFRR) profile.</p> <p>When you configure HFRR, an optional ARP prefix limit sets a maximum for the number of ARP routes and, therefore FRR routes created for each HFRR profile in the routing table. This limit prevents ARP attacks from exhausting the virtual memory on the routing devices.</p> <p>There are two configuration statements (<code>global-arp-prefix-limit</code> and <code>arp-prefix-limit</code>) that set the ARP prefix limit, one at the global <code>[edit routing-options host-fast-reroute]</code> hierarchy level and the other at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, respectively. The global <code>global-arp-prefix-limit</code> statement sets a default ARP prefix limit for all HFRR profiles configured on the routing device. The <code>arp-prefix-limit</code> statement overrides the <code>global-arp-prefix-limit</code> for that HFRR profile for that protected interface.</p> <p>Warning system log messages begin when the ARP routes in an HFRR profile reaches 80% of the configured limit. When the number crosses the 100% threshold, the HFRR profile is deactivated. When this happens, all ARP/FRR routes are deleted from the routing table. FRR routes are deleted from forwarding table as well.</p> <p>After the HFRR profile is deactivated, a blackout timer is started. The timeout value of this timer is the ARP cache timeout (kernel timeout) + the supplementary blackout timer.</p> <p>There are global and per-HFRR CLI statements (<code>global-supplementary-blackout-timer</code> and <code>supplementary-blackout-timer</code>) to configure the supplementary blackout timer. The global value is at the <code>[edit routing-options host-fast-reroute]</code> hierarchy level and applies to all HFRR profiles on the routing device. The value for the routing-instance interface is at the <code>[edit routing-instances <i>instance-name</i> routing-options interface <i>interface-name</i>]</code> hierarchy level, and overrides the global value for that HFRR profile only.</p> <p>When the blackout timer expires, the HFRR profile is reactivated, and the Junos OS relearns the ARP routes and re-creates the HFRR routes. If the ARP prefix limit is not exceeded again, the HFRR routes will be up.</p> <p>If an HFRR profile is in the deactivated state, a reevaluation of the ARP state is preformed during every commit operation or whenever the routing process (rpd) is restarted with the <code>restart routing</code> command.</p>

<b>Default</b>	If you omit the <b>supplementary-blackout-timer</b> statement, the <b>global-supplementary-blackout-timer</b> takes effect for all HFRR profiles on the device. If you omit both of these statements, the blackout timeout value is set to the ARP cache timeout value, which by default is 20 minutes and is configurable with the <b>aging-timer</b> statement at the <b>[edit system arp]</b> hierarchy level.
<b>Options</b>	<b>minutes</b> —Number of minutes, in addition to the ARP cache timeout value, before an HFRR profile is reactivated after the ARP prefix limit is exceeded. <b>Range:</b> 1 through 15 minutes
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link Protection with Host Fast Reroute on page 911</a></li></ul>

---

## target (Routing Instances MVPN)

---

<b>Syntax</b>	<pre>target target-value {     receiver target-value;     sender target-value; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	Specify the target value when importing sender and receiver site routes.
<b>Options</b>	<b>target-value</b> —Specify the target value when importing sender and receiver site routes. <b>receiver</b> —Specify the target community used when importing receiver site routes. <b>sender</b> —Specify the target community used when importing sender site routes.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Import Target Receiver and Sender for an MBGP MVPN on page 589</a></li></ul>



## template(Protocols MVPN)

```

Syntax  template template-name {
        all-regions {
            incoming;
            ingress-replication {
                create-new-ucast-tunnel;
                label-switched-path {
                    label-switched-path-template (Multicast) {
                        (default-template | lsp-template-name);
                    }
                }
            }
        }
        ldp-p2mp;
        rsvp-te {
            label-switched-path-template (Multicast) {
                (default-template | lsp-template-name);
            }
            static-lsp static-lsp;
        }
    }
    region region-name{
        incoming;
        ingress-replication {
            create-new-ucast-tunnel;
            label-switched-path {
                label-switched-path-template (Multicast){
                    (default-template | lsp-template-name);
                }
            }
        }
        ldp-p2mp;
        rsvp-te {
            label-switched-path-template (Multicast) {
                (default-template | lsp-template-name);
            }
            static-lsp static-lsp;
        }
    }
}

```

**Hierarchy Level** [edit protocols mvpn inter-region-template]

**Release Information** Statement introduced in Junos OS Release 15.1.

**Description** A template is configured on a transit area boundary router (ABR) to define the tunnel type for each region. If a template is not configured on the ABR then the ABR simply reflects the routes without performing any inter-region segmentation.

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**

- [Segmented Inter-Area Point-to-Multipoint Label-Switched Paths Overview on page 471](#)
- [Example: Configuring Segmented Inter-Area P2MP LSP on page 478](#)
- [Configuring Segmented Inter-Area P2MP LSP on page 476](#)
- [all-regions on page 1042](#)
- [inter-region on page 1082](#)
- [inter-region-template on page 1085](#)
- [inter-region-segmented on page 1083](#)
- [region on page 1122](#)

## threshold-rate

<b>Syntax</b>	<code>threshold-rate kbps;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i>]</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> wildcard-source]</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Specify the data threshold required before a new tunnel is created for a dynamic selective point-to-multipoint LSP. This statement is part of the configuration for point-to-multipoint LSPs for MBGP MVPNs and PIM-SSM GRE or RSVP-TE selective provider tunnels.
<b>Options</b>	<p><b>number</b>—Specify the data threshold required before a new tunnel is created.</p> <p><b>Range:</b> 0 through 1,000,000 kilobits per second. Specifying 0 is equivalent to not including the statement.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 469</a></li> <li>• <a href="#">Configuring PIM-SSM GRE Selective Provider Tunnels on page 597</a></li> </ul>

## traceoptions (Protocols MVPN)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvpn</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvpn</a> ], [edit protocols <a href="#">mvpn</a> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvpn</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support at the [edit protocols <a href="#">mvpn</a> ] hierarchy level introduced in Junos OS Release 13.3.
<b>Description</b>	Trace traffic flowing through a Multicast BGP (MBGP) MVPN.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches this size, it is renamed <b><i>trace-file.0</i></b>. When <b><i>trace-file</i></b> again reaches its maximum size, <b><i>trace-file.0</i></b> is renamed <b><i>trace-file.1</i></b> and <b><i>trace-file</i></b> is renamed <b><i>trace-file.0</i></b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can specify any of the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All multicast VPN tracing options</li><li>• <b>cmcast-join</b>—Multicast VPN C-multicast join routes</li><li>• <b>error</b>—Error conditions</li><li>• <b>general</b>—General events</li><li>• <b>inter-as-ad</b>—Multicast VPN inter-AS automatic discovery routes</li><li>• <b>intra-as-ad</b>—Multicast VPN intra-AS automatic discovery routes</li><li>• <b>leaf-ad</b>—Multicast VPN leaf automatic discovery routes</li></ul>

- **mdt-safi-ad**—Multicast VPN MDT SAFI automatic discovery routes
- **nlri**—Multicast VPN advertisements received or sent by means of the BGP
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **source-active**—Multicast VPN source active routes
- **spmsi-ad**—Multicast VPN SPMSI auto discovery active routes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing
- **tunnel**—Provider tunnel events
- **umh**—Upstream multicast hop (UMH) events

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing flag
- **receive**—Trace received packets
- **send**—Trace sent packets

**no-world-readable**—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 **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 also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—Allow any user to read the log file.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

- Related Documentation**
- [Tracing MBGP MVPN Traffic and Operations on page 633](#)

## traffic-statistics (Protocols BGP)

<b>Syntax</b>	<pre>traffic-statistics {     labeled-path     file <i>filename</i> &lt;world-readable   no-world-readable&gt;;     interval <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet   inet6) labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet   inet6) labeled-unicast],</p> <p>[edit protocols bgp family (inet   inet6) labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet   inet6) labeled-unicast]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.</p> <p><b>labeled-path</b> introduced in Junos OS Release 18.1R1 for the MX Series.</p>
<b>Description</b>	Enable the collection of traffic statistics for interprovider or carrier-of-carriers VPNs.
<b>Options</b>	<p><b>file <i>filename</i></b>—Specify a filename for the BGP labeled-unicast traffic statistics file. If you do not specify a filename, statistics are still collected but can only be viewed by using the <b>show bgp group traffic statistics <i>group-name</i></b> command.</p> <p><b>interval <i>seconds</i></b>—Specify how often BGP labeled-unicast traffic statistics are collected.</p> <p><b>labeled-path</b>—Specify this option to collect labeled path traffic statistics of ingress BGP nodes.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics on page 314</a></li> <li>• <a href="#">MPLS Feature Support on QFX Series and EX4600 Switches</a></li> <li>• <a href="#">Understanding Interprovider and Carrier-of-Carriers VPNs on page 312</a></li> </ul>

## tunnel-limit (Routing Instances Provider Tunnel Selective)

<b>Syntax</b>	tunnel-limit <i>number</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Specify a limit on the number of selective tunnels that can be created for an LSP. This limit can be applied to the following types of selective tunnels: <ul style="list-style-type: none"> <li>• Ingress replication tunnels</li> <li>• LDP-signaled LSP</li> <li>• LDP point-to-multipoint LSP</li> <li>• PIM-SSM provider tunnel</li> <li>• RSVP-signaled LSP</li> <li>• RSVP-signaled point-to-multipoint LSP</li> </ul>
<b>Options</b>	<i>number</i> —Specify the tunnel limit. <b>Range:</b> 0 through 1024
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 470</a></li> <li>• <a href="#">selective on page 1129</a></li> <li>• <a href="#">wildcard-source on page 1158</a></li> </ul>

## unicast (Route Target Community)

---

<b>Syntax</b>	<pre>unicast {     receiver;     sender; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	Specify the same target community configured for unicast.
<b>Options</b>	<b>receiver</b> —Specify the unicast target community used when importing receiver site routes. <b>sender</b> —Specify the unicast target community used when importing sender site routes.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Import Target Unicast Parameters for an MBGP MVPN on page 590</a></li></ul>



## unicast (Virtual Tunnel in Routing Instances)

<b>Syntax</b>	unicast;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i> ], [edit routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4.
<b>Description</b>	In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used for unicast traffic only.
<b>Default</b>	If you omit this statement, the VT interface can be used for both multicast and unicast traffic.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 600</a></li> <li>• <a href="#">Example: Configuring MBGP MVPN Extranets</a></li> </ul>

## vpn-localization

---

<b>Syntax</b>	<pre>vpn-localization {   vpn-core-facing-default;   vpn-core-facing-only; }</pre>
<b>Hierarchy Level</b>	[edit chassis fpc <i>fpc-slot-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2.
<b>Description</b>	Configure FPC as either core-facing default or core-facing only in order to localize the routes.
<b>Options</b>	<p><b>vpn-core-facing-default</b>— The core-facing FPC installs all the routes and next hops of the CE-facing interfaces.</p> <p><b>vpn-core-facing-only</b>— The core-facing FPC installs all the routes and does not store next hops of the CE-facing interfaces.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding VRF Localization in Layer 3 VPNs on page 145</a></li><li>• <a href="#">Example: Improving Scalability Using VRF Localization for Layer 3 VPNs on page 147</a></li><li>• <a href="#">localized-fib on page 1095</a></li></ul>

## vpn-unequal-cost

<b>Syntax</b>	vpn-unequal-cost { equal-external-internal; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multipath], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> multipath], [edit routing-instances <i>routing-instance-name</i> routing-options multipath], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> multipath]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. The <b>equal-external-internal</b> option was added for Junos OS Release 8.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	When the <b>vpn-unequal-cost</b> statement is configured, protocol-independent load balancing is applied to VPN routes that are equal to interior gateway protocol (IGP) metrics. If the <b>vpn-unequal-cost</b> statement is not configured, then protocol-independent load balancing is applied to VPN routes that are equal to the router identifier.
<b>Options</b>	<b>equal-external-internal</b> —Specifies that both external and internal BGP paths can be selected for multipath.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Load Balancing for Layer 3 VPNs on page 787</a></li> <li>• <a href="#">Load Balancing and IP Header Filtering for Layer 3 VPNs on page 770</a></li> </ul>

## vrf-advertise-selective

---

<b>Syntax</b>	<pre>vrf-advertise-selective {     family {         inet-mvpn;         inet6-mvpn;     } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	<p>Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.</p> <p>If you configure the <b>vrf-advertise-selective</b> statement without any of its options, the router or switch has the same behavior as if you configured the <b>no-vrf-advertise</b> statement. All VPN routes are prevented from being advertised from a VRF routing instance to the remote PE routers. This behavior is useful for hub-and-spoke configurations, enabling you to configure a PE router to not advertise VPN routes from the primary (hub) instance. Instead, these routes are advertised from the secondary (downstream) instance.</p> <p>The options are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Limiting Routes to Be Advertised by an MVPN VRF Instance on page 590</a></li><li>• <a href="#">no-vrf-advertise on page 1107</a></li></ul>

## vrf-propagate-ttl

---

<b>Syntax</b>	vrf-propagate-ttl;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Enable normal time-to-live (TTL) decrementing in a VRF routing instance. You configure this statement once per routing instance, and it affects only RSVP-signaled or LDP-signaled LSPs in the routing instance. When this router acts as an ingress router for an LSP, it pushes an MPLS copies the TTL value from the IP packet header. When the router acts as the penultimate router, it pops the MPLS header and writes the MPLS TTL into the IP packet.
<b>Default</b>	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP. This statement explicitly configures the default behavior for the VRF routing instance and is useful for overriding the <b>no-propagate-ttl</b> configured globally on the router at the <b>[edit protocols mpls]</b> hierarchy level.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing on page 989</a></li> <li>• <i>Disabling Normal TTL Decrementing</i></li> </ul>

## vrf-table-label

<b>Syntax</b>	<pre>vrf-table-label {     source-class-usage;     static; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the <b>source-class-usage</b> statement added in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 15.1F5 and 16.1R2 for PTX5000 routers with third-generation FPCs installed.</p> <p>Statement introduced in Junos OS Release 15.1F6 and 16.1R2 for PTX3000 routers with third-generation FPCs.</p> <p>Statement introduced in Junos OS Release 16.1X65 and 17.2R1 for PTX1000 routers.</p> <p>Support for the <b>static</b> statement added in Junos OS Release 17.2.</p>
<b>Description</b>	<p>Map the inner label of a packet to a specific VPN routing and forwarding (VRF) instance. This allows the examination of the encapsulated IP header. The first lookup is done on the VPN label to determine which VRF instance to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts.</p> <p>When you include the <b>vrf-table-label</b> statement in the configuration of a VRF routing instance, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table. Any routes in the VRF routing table are advertised with the LSI logical interface label allocated for the VRF routing table. When packets destined for the VRF routing instance arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table.</p> <p>All routes in a VRF routing instance configured with this option are advertised with one label allocated per VRF.</p>



### NOTE:

- The **vrf-table-label** statement is supported on PTX5000 and PTX3000 routers only when third-generation FPCs are installed on the router and **enhanced-ip** command is configured on the chassis.
- Starting in Junos OS Release 17.2, you can configure the **enhanced-ip** command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability. The **vrf-table-label** space is increased to at least 16,000, if the platform can support the scale.

**Options** The remaining statements are explained separately.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- [Filtering Packets in Layer 3 VPNs Based on IP Headers on page 159](#)
- [Configuring EXP-Based Traffic Classification for VPLS](#)
- [Load Balancing and IP Header Filtering for Layer 3 VPNs on page 770](#)

## wildcard-group-inet

<b>Syntax</b>	<pre> wildcard-group-inet {   wildcard-source {     inter-region-segmented{       fan-out <i>fan-out value</i>;     }     ldp-p2mp;     pim-ssm {       group-range <i>multicast-prefix</i>;     }     rsvp-te {       label-switched-path-template {         (default-template   <i>lsp-template-name</i>);       }       static-lsp <i>lsp-name</i>;     }     threshold-rate <i>number</i>;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>The <b>inter-region-segmented</b> statement added in Junos OS Release 15.1.</p>
<b>Description</b>	<p>Configure a wildcard group matching any group IPv4 address.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">wildcard-group-inet6 on page 1157</a></li> <li>• <a href="#">Example: Configuring Selective Provider Tunnels Using Wildcards on page 565</a></li> <li>• <a href="#">Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559</a></li> <li>• <a href="#">Configuring a Selective Provider Tunnel Using Wildcards on page 564</a></li> </ul>



## wildcard-group-inet6

<b>Syntax</b>	<pre> wildcard-group-inet6 {   wildcard-source {     inter-region-segmented{       fan-out <i>fan-out value</i>;     }     ldp-p2mp;     pim-ssm {       group-range <i>multicast-prefix</i>;     }     rsvp-te {       label-switched-path-template {         (default-template   <i>lsp-template-name</i>);       }       static-lsp <i>lsp-name</i>;     }     threshold-rate <i>number</i>;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>The <b>inter-region-segmented</b> statement added in Junos OS Release 15.1.</p>
<b>Description</b>	<p>Configure a wildcard group matching any group IPv6 address.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">wildcard-group-inet on page 1156</a></li> <li>• <a href="#">Example: Configuring Selective Provider Tunnels Using Wildcards on page 565</a></li> <li>• <a href="#">Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559</a></li> <li>• <a href="#">Configuring a Selective Provider Tunnel Using Wildcards on page 564</a></li> </ul>

## wildcard-source (Selective Provider Tunnels)

<b>Syntax</b>	<pre>wildcard-source {   inter-region-segmented {     fan-out fan-out value;   }   ldp-p2mp;   pim-ssm {     group-range multicast-prefix;   }   rsvp-te {     label-switched-path-template {       (default-template   lsp-template-name);     }   }   static-lsp lsp-name; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems logical-system-name routing-instances routing-instance-name   provider-tunnel selective group group-prefix], [edit logical-systems logical-system-name routing-instances routing-instance-name   provider-tunnel selective wildcard-group-inet], [edit logical-systems logical-system-name routing-instances routing-instance-name   provider-tunnel selective wildcard-group-inet6], [edit routing-instances routing-instance-name provider-tunnel selective group group-prefix], [edit logical-systems logical-system-name routing-instances routing-instance-name   provider-tunnel selective wildcard-group-inet], [edit logical-systems logical-system-name routing-instances routing-instance-name   provider-tunnel selective wildcard-group-inet6], [edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet], [edit routing-instances routing-instance-name provider-tunnel selective wildcard-group-inet6]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>The <b>inter-region-segmented</b> statement added in Junos OS Release 15.1.</p>
<b>Description</b>	<p>Configure a selective provider tunnel for a shared tree using a wildcard source.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">wildcard-group-inet on page 1156</a></li> <li>• <a href="#">wildcard-group-inet6 on page 1157</a></li> <li>• <a href="#">Example: Configuring Selective Provider Tunnels Using Wildcards on page 565</a></li> <li>• <a href="#">Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 559</a></li> <li>• <a href="#">Configuring a Selective Provider Tunnel Using Wildcards on page 564</a></li> </ul>

## CHAPTER 51

# Operational Commands

- ping mpls l3vpn
- show hfr profiles
- show ingress-replication mvpn
- show mvpn c-multicast
- show mvpn instance
- show mvpn neighbor
- show route vpn-localization

## ping mpls l3vpn

---

**Syntax** ping mpls l3vpn prefix *prefix-name*  
<*l3vpn-name*>  
<bottom-label-ttl>  
<count *count*>  
<destination *address*>  
<detail>  
<exp *forwarding-class*>  
<logical-system (all | *logical-system-name*)>  
<size *bytes*>  
<source *source-address*>  
<sweep>

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

**Description** Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a **ping mpls l3vpn** command.

**Options** **bottom-label-ttl**—(Optional) Display the time-to-live value for the bottom label in the label stack.

**count *count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination *address***—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

***l3vpn-name***—(Optional) Layer 3 VPN name.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.

**prefix *prefix-name***—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.

**size *bytes***—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size

of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

<b>Additional Information</b>	<p>You must configure MPLS at the <b>[edit protocols mpls]</b> hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.</p> <p>In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.</p> <p>If the Layer 3 VPN traffic transits a route reflector within the network, the <b>ping mpls l3vpn</b> command does not work.</p>
<b>Required Privilege Level</b>	network
<b>List of Sample Output</b>	<p><a href="#">ping mpls l3vpn on page 1161</a>  <a href="#">ping mpls l3vpn detail on page 1161</a></p>
<b>Output Fields</b>	<p>When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.</p>

## Sample Output

### ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!
--- 1sping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

### ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
```

```
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## show hfr profiles

<b>Syntax</b>	show hfr profiles <brief   extensive>
<b>Release Information</b>	Command introduced in Junos OS Release 12.2.
<b>Description</b>	<p>Display host fast reroute (HFRR) profile information.</p> <p>HFRR adds a precomputed protection path into the Packet Forwarding Engine, such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the Packet Forwarding Engine can use another path without having to wait for the router or the protocols to provide updated forwarding information.</p>
<b>Options</b>	<p><b>none</b>—Display information about HRFF profiles.</p> <p><b>brief   extensive</b>—(Optional) Display the specified level of output.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Link Protection with Host Fast Reroute on page 911</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show hfr profiles on page 1164</a>
<b>Output Fields</b>	<p><a href="#">Table 15 on page 1163</a> describes the output fields for the <b>show hfr profiles</b> command. Output fields are listed in the approximate order in which they appear.</p>

*Table 15: show hfr profiles Output Fields*

Field Name	Field Description
HFRR	
HFRR current state	Status of the HFRR profile: HFRR_ACTIVE, HFRR_INACTIVE, HFRR_IFLH-NOT-CONF, and so on.
HFRR Prefix limit blackout timer expiry (in secs)	Time interval between an HFRR profile becoming inactive on exceeding the ARP prefix limit, and the profile starting the SYNC process.
HFRR prefix limit hit count	Number of times that an HFRR profile becomes inactive on exceeding the ARP prefix limit.
HFRR protected IFL name	Interface configured for the HFRR feature.
HFRR protected IFL handle	
HFRR routing instance name	The routing instance in which the HFRR interface is configured.

Table 15: show hfr profiles Output Fields (continued)

Field Name	Field Description
HFRR routing instance handle	
HFRR sync BG scheduled	
HFRR RTS filter on	
HFRR delete BG scheduled	
HFRR ARP prefix limit	Configured ARP prefix limit.
HFRR ARP supplementary blackout timeout (in mins)	Supplementary time-out value configured for profile to be inactive when it hits ARP prefix limit.
HFRR number of ARP routes learned	Number of ARP routes learned on the configured interface.
HFRR number of FRR routes created	Number of ARP routes created on the configured interface.

## Sample Output

### show hfr profiles

```

user@host> show hfr profiles
HFRR pointer: 0x9254000
HFRR current state: HFRR_ACTIVE
HFRR Prefix limit blackout timer expiry (in secs): 0
HFRR prefix limit hit count: 0
HFRR protected IFL name: ge-4/1/0.0
HFRR protected IFL handle: 0x9248738
HFRR routing instance name: test
HFRR routing instance handle: 0x9145740
HFRR sync BG scheduled: NO
HFRR RTS filter on: YES
HFRR delete BG scheduled: NO
HFRR ARP prefix limit: 0
HFRR ARP supplementary blackout timeout (in mins): 1
HFRR number of ARP routes learned: 4
HFRR number of FRR routes created: 2

```



## show ingress-replication mvpn

**Syntax** show ingress-replication mvpn

**Release Information** Command introduced in Junos OS Release 10.4.

**Description** Display the state and configuration of the ingress replication tunnels created for the MVPN application when using the **mpls-internet-multicast** routing instance type.

**Required Privilege Level** View

**List of Sample Output** [show ingress-replication mvpn on page 1165](#)

**Output Fields** [Table 16 on page 1165](#) lists the output fields for the **show ingress-replication mvpn** command. Output fields are listed in the approximate order in which they appear.

*Table 16: show ingress-replication mvpn Output Fields*

Field Name	Field Description
Ingress tunnel	Identifies the MVPN ingress replication tunnel.
Application	Identifies the application (MVPN).
Unicast tunnels	List of unicast tunnels in use.
Leaf address	Address of the tunnel.
Tunnel type	Identifies the unicast tunnel type.
Mode	Indicates whether the tunnel was created as a new tunnel for the ingress replication, or if an existing tunnel was used.
State	Indicates whether the tunnel is Up or Down.

## Sample Output

### show ingress-replication mvpn

```

user@host> show ingress-replication mvpn
Ingress Tunnel: mvpn:1
  Application: MVPN
  Unicast tunnels
    Leaf Address      Tunnel-type      Mode      State
    10.255.245.2      P2P LSP         New       Up
    10.255.245.4      P2P LSP         New       Up
Ingress Tunnel: mvpn:2
  Application: MVPN
  Unicast tunnels

```

Leaf Address	Tunnel-type	Mode	State
10.255.245.2	P2P LSP	Existing	Up

## show mvpn c-multicast

<b>Syntax</b>	show mvpn c-multicast <extensive   summary> <instance-name <i>instance-name</i> > <source-pe>
<b>Release Information</b>	Command introduced in Junos OS Release 8.4. Option to show <b>source-pe</b> introduced in Junos OS Release 15.1.
<b>Description</b>	Display the multicast VPN customer multicast route information.
<b>Options</b>	<b>extensive   summary</b> —(Optional) Display the specified level of output.  <b>instance-name <i>instance-name</i></b> —(Optional) Display output for the specified routing instance.  <b>source-pe</b> —(Optional) Display source-pe output for the specified c-multicast entries.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mvpn c-multicast on page 1168</a> <a href="#">show mvpn c-multicast summary on page 1168</a> <a href="#">show mvpn c-multicast extensive on page 1168</a> <a href="#">show mvpn c-multicast source-pe on page 1169</a>
<b>Output Fields</b>	<a href="#">Table 17 on page 1167</a> lists the output fields for the <b>show mvpn c-multicast</b> command. Output fields are listed in the approximate order in which they appear.

**Table 17: show mvpn c-multicast Output Fields**

Field Name	Field Description	Level of Output
Instance	Name of the VPN routing instance.	summary extensive none
C-mcast IPv4 (S:G)	Customer router IPv4 multicast address.	extensive none
Ptnl	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none
St	State: <ul style="list-style-type: none"> <li>DS—Represents (S,G) and is created due to (*,G)</li> <li>RM—Remote VPN route learned from the remote PE router</li> <li>St display blank—SSM group join</li> </ul>	extensive none
MVPN instance	Name of the multicast VPN routing instance	extensive none

Table 17: *show mvpn c-multicast Output Fields (continued)*

Field Name	Field Description	Level of Output
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

## Sample Output

### show mvpn c-multicast

```

user@host> show mvpn c-multicast
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.78/32:203.0.113.1/24 PIM-SM:10.255.14.144, 198.51.100.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.94/32:203.0.113.0/24 PIM-SM:10.255.14.144, 198.51.100.2      RM

```

### show mvpn c-multicast summary

```

user@host> show mvpn c-multicast summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  C-multicast IPv6 route count: 1

```

### show mvpn c-multicast extensive

```

user@host> show mvpn c-multicast extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)          Ptnl          St

```

```

192.168.195.78/32:203.0.113.1/24 PIM-SM:10.255.14.144, 198.51.100.1      RM
MVPN instance:

```

Legend for provider tunnel

I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (\*, c-g) RM -- remote VPN route

Instance: VPN-B

```

C-mcast IPv4 (S:G)          Ptnl          St
192.168.195.94/32:203.0.113.0/24 PIM-SM:10.255.14.144, 198.51.100.2      RM

```

### show mvpn c-multicast source-pe

```

user@host> show mvpn c-multicast source-pe

```

Family : INET

Family : INET6

Instance : mvpn1

MVPN Mode : RPT-SPT

C-Multicast route address: ::0:ff05::1/128

MVPN Source-PE1:

extended-community: no-advertise target:10.1.0.0:9

Route Distinguisher: 10.1.0.0:1

Autonomous system number: 1

Interface: ge-0/0/9.1 Index: 343

PIM Source-PE1:

extended-community: target:10.1.0.0:9

Route Distinguisher: 10.1.0.0:1

Autonomous system number: 1

Interface: ge-0/0/9.1 Index: 343

## show mvpn instance

<b>Syntax</b>	<pre>show mvpn instance &lt;instance-name&gt; &lt;display-tunnel-name&gt; &lt;extensive   summary&gt; &lt;inet   inet6&gt; &lt;logical-system&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.4.</p> <p>Additional details in output for extensive option introduced in Junos OS Release 15.1.</p>
<b>Description</b>	Display the multicast VPN routing instance information according the options specified.
<b>Options</b>	<p><b>instance-name</b>—(Optional) Display statistics for the specified routing instance, or press Enter without specifying an instance name to show output for all instances.</p> <p><b>display-tunnel-name</b>—(Optional) Display the ingress provider tunnel name rather than the attribute.</p> <p><b>extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display output for the specified IP type.</p> <p><b>inet   inet6</b>—(Optional) Display output for the specified IP type.</p> <p><b>logical-system</b>—(Optional) Display details for the specified logical system, or type “all”.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show mvpn instance on page 1171</a></p> <p><a href="#">show mvpn instance summary on page 1172</a></p> <p><a href="#">show mvpn instance extensive on page 1172</a></p> <p><a href="#">show mvpn instance summary (IPv6) on page 1173</a></p>
<b>Output Fields</b>	<p><a href="#">Table 18 on page 1170</a> lists the output fields for the <b>show mvpn instance</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 18: show mvpn instance Output Fields**

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	extensive none
Instance	Name of the VPN routing instance.	summary extensive none
Provider tunnel	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none

Table 18: show mvpn instance Output Fields (continued)

Field Name	Field Description	Level of Output
Neighbor	Address, type of provider tunnel (I-P-tnl, inclusive provider tunnel and S-P-tnl, selective provider tunnel) and provider tunnel for each neighbor.	extensive none
C-mcast IPv4 (S:G)	Customer IPv4 router multicast address.	extensive none
C-mcast IPv6 (S:G)	Customer IPv6 router multicast address.	extensive none
Ptnl	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none
St	State: <ul style="list-style-type: none"> <li>DS—Represents (S,G) and is created due to (*G)</li> <li>RM—Remote VPN route learned from the remote PE router</li> <li>St display blank—SSM group join</li> </ul>	extensive none
Neighbor count	Number of neighbors associated with the multicast VPN routing instance.	summary
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

## Sample Output

### show mvpn instance

```

user@host> show mvpn instance
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 198.51.100.1
  Neighbor                    I-P-tnl
10.255.14.160                  PIM-SM:10.255.14.160, 198.51.100.1
10.255.70.17                   PIM-SM:10.255.70.17, 198.51.100.1
  C-mcast IPv4 (S:G)          Ptnl                      St
192.168.195.78/32:203.0.113.0/24 PIM-SM:10.255.14.144, 198.51.100.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-B

```

```

Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 198.51.100.2
Neighbor                               I-P-tnl
10.255.14.160                          PIM-SM:10.255.14.160, 198.51.100.2
10.255.70.17                           PIM-SM:10.255.70.17, 198.51.100.2
C-mcast IPv4 (S:G)                     Ptnl                      St
192.168.195.94/32:203.0.113.1/24 PIM-SM:10.255.14.144, 198.51.100.2      RM

```

## Sample Output

### show mvpn instance summary

```
user@host> show mvpn instance summary
```

```
MVPN Summary:
```

```
Family: INET
```

```
Family: INET6
```

```
Instance: mvpn1
```

```
Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
```

```
Hot Root Standby: Disabled. Reason: Not enabled by configuration.
```

```
Neighbor count: 3
```

```
C-multicast IPv6 route count: 1
```

## Sample Output

### show mvpn instance extensive

```
user@host> show mvpn instance extensive
```

```
MVPN instance:
```

```
Family : INET
```

```
Instance : vpn_blue
```

```
Customer Source: 10.1.1.1
```

```
RT-Import Target: 192.168.1.1:100
```

```
Route-Distinguisher: 192.168.1.1:100
```

```
Source-AS: 65000
```

```
Via unicast route: 10.1.0.0/16 in vpn-blue.inet.0
```

```
Candidate Source PE Set:
```

```
RT-Import 192.168.1.1:100, RD 1111:22222, Source-AS 65000
```

```
RT-Import 192.168.2.2:100, RD 1111:22222, Source-AS 65000
```

```
RT-Import 192.168.3.3:100, RD 1111:22222, Source-AS 65000
```

‘Extensive’ output will show everything in ‘detail’ output and add the list of bound c-multicast routes.

```
> show mvpn source 10.1.1.1 instance vpn_blue extensive
```

```
Family : INET
```

```
Instance : vpn_blue
```

```
Customer Source: 10.1.1.1
```

```
RT-Import Target: 192.168.1.1:100
```

```
Route-Distinguisher: 192.168.1.1:100
```

```
Source-AS: 65000
```

```
Via unicast route: 10.1.0.0/16 in vpn-blue.inet.0
```

```
Candidate Source PE Set:
```

```
RT-Import 192.168.1.1:100, RD 1111:22222, Source-AS 65000
```

```
RT-Import 192.168.2.2:100, RD 1111:22222, Source-AS 65000
```

```
RT-Import 192.168.3.3:100, RD 1111:22222, Source-AS 65000
```

```
Customer-Multicast Routes:
```



```
10.1.1.1/32:198.51.100.3/24
10.1.1.1/32:198.51.100.3/24
```

#### show mvpn instance summary (IPv6)

```
user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
  C-multicast IPv6 route count: 2
Instance: VPN-B
  C-multicast IPv6 route count: 2
```

## show mvpn neighbor

<b>Syntax</b>	show mvpn neighbor <extensive   summary> <inet   inet6> <instance <i>instance-name</i>   neighbor-address <i>address</i> > <logical-system <i>logical-system-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.4.
<b>Description</b>	Display multicast VPN neighbor information.
<b>Options</b>	<p><b>extensive   summary</b>—(Optional) Display the specified level of output for all multicast VPN neighbors.</p> <p><b>inet   inet6</b>—(Optional) Display IPv4 or IPv6 information for all multicast VPN neighbors.</p> <p><b>instance <i>instance-name</i>   neighbor-address <i>address</i></b>—(Optional) Display multicast VPN neighbor information for the specified instance or the specified neighbor.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display multicast VPN neighbor information for the specified logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mvpn neighbor on page 1175</a> <a href="#">show mvpn neighbor extensive on page 1175</a> <a href="#">show mvpn neighbor extensive on page 1176</a> <a href="#">show mvpn neighbor instance-name on page 1176</a> <a href="#">show mvpn neighbor neighbor-address on page 1176</a> <a href="#">show mvpn neighbor neighbor-address summary on page 1177</a> <a href="#">show mvpn neighbor neighbor-address extensive on page 1177</a> <a href="#">show mvpn neighbor neighbor-address instance-name on page 1177</a> <a href="#">show mvpn neighbor summary on page 1178</a>
<b>Output Fields</b>	<a href="#">Table 19 on page 1174</a> lists the output fields for the <b>show mvpn neighbor</b> command. Output fields are listed in the approximate order in which they appear.

Table 19: show mvpn neighbor Output Fields

Field Name	Field Description	Level of Output
<b>MVPN instance</b>	Name of the multicast VPN routing instance	extensive none
<b>Instance</b>	Name of the VPN routing instance.	summary extensive none

Table 19: show mvpn neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
Neighbor	Address, type of provider tunnel (I-P-tnl, inclusive provider tunnel and S-P-tnl, selective provider tunnel) and provider tunnel for each neighbor.	extensive none
Provider tunnel	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none

## Sample Output

### show mvpn neighbor

```

user@host> show mvpn neighbor
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-A
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 192.0.2.1
10.255.70.17                     PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-B
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 192.0.2.2
10.255.70.17                     PIM-SM:10.255.70.17, 192.0.2.2

```

## Sample Output

### show mvpn neighbor extensive

```

user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-A
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 192.0.2.1
10.255.70.17                     PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

```

```

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                               I-P-tnl
10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.2
10.255.70.17                         PIM-SM:10.255.70.17, 192.0.2.2

```

### show mvpn neighbor extensive

```

user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: mvpn-a
Neighbor                               I-P-tnl
10.255.72.45                           LDP P2MP:10.255.72.50, lsp-id 1
10.255.72.50

```

## Sample Output

### show mvpn neighbor instance-name

```

user@host> show mvpn neighbor instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                               I-P-tnl
10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
10.255.70.17                         PIM-SM:10.255.70.17, 192.0.2.1

```

## Sample Output

### show mvpn neighbor neighbor-address

```

user@host> show mvpn neighbor neighbor-address 10.255.14.160
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                               I-P-tnl
10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

```

```

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                             I-P-tnl
10.255.14.160                        PIM-SM:10.255.14.160, 192.0.2.2

```

## Sample Output

### show mvpn neighbor neighbor-address summary

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 summary
MVPN Summary:
Instance: VPN-A
Instance: VPN-B

```

## Sample Output

### show mvpn neighbor neighbor-address extensive

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl
10.255.70.17                        PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                             I-P-tnl
10.255.70.17                        PIM-SM:10.255.70.17, 192.0.2.2

```

## Sample Output

### show mvpn neighbor neighbor-address instance-name

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl
10.255.70.17                        PIM-SM:10.255.70.17, 192.0.2.1

```

## Sample Output

### show mvpn neighbor summary

```
user@host> show mvpn neighbor summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  Neighbor count: 3
```

## show route vpn-localization

<b>Syntax</b>	show route vpn-localization <brief   detail   terse> <family (inet   inet6)> <vpn-name <i>vpn-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 14.2.
<b>Description</b>	Display the localization information of all the VRFs in the system.
<b>Options</b>	<p><b>none</b>— Display the localization information of all the VRFs in the system.</p> <p><b>brief   detail   terse</b>—(Optional) Display the specified level of output.</p> <p><b>family (inet   inet6)</b>—(Optional) Display the localized information of all the VRFs in the system for inet or inet6 family.</p> <p><b>vpn-name <i>vpn-name</i></b>—(Optional) Display the localized information of the specific VPN.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show route vpn-localization on page 1180</a></p> <p><a href="#">show route vpn-localization vpn-name family inet on page 1182</a></p> <p><a href="#">show route vpn-localization vpn-name family inet detail on page 1182</a></p> <p><a href="#">show route vpn-localization terse on page 1182</a></p>
<b>Output Fields</b>	<a href="#">Table 20 on page 1179</a> describes the output fields for the <b>show route vpn-localization</b> command. Output fields are listed in the approximate order in which they appear.

*Table 20: show route vpn-localization Output Fields*

Field Name	Field Description	Level of Output
Routing table	Name of the routing table	none <b>detail</b>
Index	Index of the routing table.	none <b>detail</b>
Address Family	Address family of the routing table (inet/inet6).	none <b>detail</b>
Localization status	Localization status of the routing table and can be one of the following: <ul style="list-style-type: none"> <li>Complete - The routing table successfully installed on all the local FPCs.</li> <li>Pending - The localization change on this routing table is yet to be processed in the pending state. The routing table might not be present on all the local FPCs or can be present on the non-local FPCs.</li> <li>Replaying - The routing table is being installed on the local FPCs. A localization change involving the addition of new local FPCs requires the system to replay the routing table for newly added FPCs.</li> </ul>	none <b>detail</b>

Table 20: show route vpn-localization Output Fields (continued)

Field Name	Field Description	Level of Output
Local FPC's	List of FPCs on which the routing table should be installed, as per the configuration.	none
Local FPC (configured)	List of FPCs on which the routing table should be installed, as per the configuration.	detail
Local FPC (Current state)	List of FPCs on which the routing table is currently installed.	detail
Table Name	Name of the routing table.	terse
Local'd	The localization status of VRF. If VRF is localized, then the status is Yes; otherwise, it is No.	terse
Table Index	Index of the routing table.	terse
Local FPC's Configured	List of FPCs, in hexadecimal, on which the routing table should be installed as per the configuration.	terse
Local FPC's Actual	List of FPCs, in hexadecimal, on which the routing table is currently installed.	terse
State	Localization status of the routing table and can be one of the following: <ul style="list-style-type: none"> <li>Complete - The routing table successfully installed on the local FPCs.</li> <li>Pending - The localization change on this routing table is yet to be processed in the pending state. The routing table might not be present on all the local FPCs or can be present on the non-local FPCs.</li> <li>Replaying - The routing table is being installed on the local FPCs. A localization change involving addition of new local FPCs requires the system to replay the routing table for the newly added FPCs.</li> </ul>	terse

## Sample Output

### show route vpn-localization

```

user@host> show route vpn-localization

Routing table: CE-1.inet, Localized
  Index: 7, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 13 18

Routing table: CE-1.inet6, Localized
  Index: 7, Address Family: inet6, Localization status: Complete
  Local FPC's: 4 5 13 18

Routing table: CE-10.inet, Localized
  Index: 8, Address Family: inet, Localization status: Complete
  Local FPC's: 4 5 13

Routing table: CE-10.inet6, Localized
  Index: 8, Address Family: inet6, Localization status: Complete

```



Local FPC's: 4 5 13

Routing table: CE-2.inet, Localized

Index: 9, Address Family: inet, Localization status: Complete

Local FPC's: 4 5 10 13

Routing table: CE-2.inet6, Localized

Index: 9, Address Family: inet6, Localization status: Complete

Local FPC's: 4 5 10 13

Routing table: CE-3.inet, Localized

Index: 10, Address Family: inet, Localization status: Complete

Local FPC's: 4 5 9 13

Routing table: CE-3.inet6, Localized

Index: 10, Address Family: inet6, Localization status: Complete

Local FPC's: 4 5 9 13

Routing table: CE-4.inet, Localized

Index: 11, Address Family: inet, Localization status: Complete

Local FPC's: 3 4 5 13

Routing table: CE-4.inet6, Localized

Index: 11, Address Family: inet6, Localization status: Complete

Local FPC's: 3 4 5 13

Routing table: CE-5.inet, Localized

Index: 12, Address Family: inet, Localization status: Complete

Local FPC's: 4 5 13

Routing table: CE-5.inet6, Localized

Index: 12, Address Family: inet6, Localization status: Complete

Local FPC's: 4 5 13

Routing table: CE-6.inet, Localized

Index: 13, Address Family: inet, Localization status: Complete

Local FPC's: 4 5 13 18

Routing table: CE-6.inet6, Localized

Index: 13, Address Family: inet6, Localization status: Complete

Local FPC's: 4 5 13 18

Routing table: CE-7.inet, Localized

Index: 14, Address Family: inet, Localization status: Complete

Local FPC's: 4 5 10 13

Routing table: CE-7.inet6, Localized

Index: 14, Address Family: inet6, Localization status: Complete

Local FPC's: 4 5 10 13

Routing table: CE-8.inet, Localized

Index: 15, Address Family: inet, Localization status: Complete

Local FPC's: 4 5 9 13

Routing table: CE-8.inet6, Localized

Index: 15, Address Family: inet6, Localization status: Complete

Local FPC's: 4 5 9 13

Routing table: CE-9.inet, Localized

Index: 16, Address Family: inet, Localization status: Complete

Local FPC's: 3 4 5 13

```

Routing table: CE-9.inet6, Localized
  Index: 16, Address Family: inet6, Localization status: Complete
  Local FPC's: 3 4 5 13

```

#### show route vpn-localization vpn-name vpn-name family inet

```
user@host> show route vpn-localization vpn-name vpn-m200 family inet
```

```

Routing table: vpn-m200.inet, Localized
  Index: 7, Address Family: inet, Localization status: Complete
  Local FPC's: 0 1

```

#### show route vpn-localization vpn-name vpn-name family inet detail

```
user@host> show route vpn-localization vpn-name vpn-m200 family inet detail
```

```

Routing table: vpn-m200.inet, Non-localized
  Index: 7, Address Family: inet
  Localization status: Complete
  Local FPC (configured):
    List: All
  Local FPC (Current state):
    List: All

```

#### show route vpn-localization terse

```
user@host> show route vpn-localization terse
```

Table Name	Local'd	Table Index	Local FPC's Configured	Actual	State
vpn-m200.inet	Yes	7	0000000000000003	0000000000000003	Complete
vpn-m200.inet6	Yes	7	0000000000000003	0000000000000003	Complete
vpn-m300.inet	Yes	8	0000000000000005	0000000000000005	Complete
vpn-m300.inet6	Yes	8	0000000000000005	0000000000000005	Complete
vpn-m500.inet	Yes	9	0000000000000021	0000000000000021	Complete
vpn-m500.inet6	Yes	9	0000000000000021	0000000000000021	Complete