

Junos® OS

Multicast Protocols User Guide

Published
2024-12-09

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the 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 Multicast Protocols User Guide

Copyright © 2024 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://support.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 This Guide | xxvi

1

Overview

Understanding Multicast | 2

Multicast Overview | 2

Understanding Layer 3 Multicast Functionality on the SRX5K-MPC | 16

Multicast Configuration Overview | 17

IPv6 Multicast Flow | 19

Supported IP Multicast Protocol Standards | 21

Ingress and Egress Multicast Replication using Recycle Replication Model | 23

2

Managing Group Membership

Configuring IGMP and MLD | 25

Configuring IGMP | 25

Understanding Group Membership Protocols | 26

Understanding IGMP | 27

Configuring IGMP | 29

Enabling IGMP | 31

Modifying the IGMP Host-Query Message Interval | 32

Modifying the IGMP Query Response Interval | 33

Specifying Immediate-Leave Host Removal for IGMP | 34

Filtering Unwanted IGMP Reports at the IGMP Interface Level | 35

Accepting IGMP Messages from Remote Subnetworks | 36

Modifying the IGMP Last-Member Query Interval | 37

Modifying the IGMP Robustness Variable | 38

Limiting the Maximum IGMP Message Rate | 40

Changing the IGMP Version | 40

Enabling IGMP Static Group Membership | 41

Recording IGMP Join and Leave Events | 50

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces | 52

- Tracing IGMP Protocol Traffic | 54
- Disabling IGMP | 56
- IGMP and Nonstop Active Routing | 57

Verifying the IGMP Version | 57

Configuring MLD | 59

- Understanding MLD | 59
- Configuring MLD | 63
- Enabling MLD | 64
- Modifying the MLD Version | 65
- Modifying the MLD Host-Query Message Interval | 66
- Modifying the MLD Query Response Interval | 67
- Modifying the MLD Last-Member Query Interval | 68
- Specifying Immediate-Leave Host Removal for MLD | 69
- Filtering Unwanted MLD Reports at the MLD Interface Level | 70
- Example: Modifying the MLD Robustness Variable | 71

- Requirements | 71
- Overview | 72
- Configuration | 72
- Verification | 73

Limiting the Maximum MLD Message Rate | 73

Enabling MLD Static Group Membership | 74

- Create a MLD Static Group Member | 74
- Automatically create static groups | 75
- Automatically increment group addresses | 77
- Specify multicast source address (in SSM mode) | 78
- Automatically specify multicast sources | 79
- Automatically increment source addresses | 81
- Exclude multicast source addresses (in SSM mode) | 82

Example: Recording MLD Join and Leave Events | 84

- Requirements | 84
- Overview | 84
- Configuration | 85
- Verification | 87

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces | 87

Disabling MLD | 89

Understanding Distributed IGMP | 90

Enabling Distributed IGMP | 92

Enabling Distributed IGMP on Static Interfaces | 92

Enabling Distributed IGMP on Dynamic Interfaces | 93

Configuring Multicast Traffic for Distributed IGMP | 93

Configuring IGMP Snooping | 96

IGMP Snooping Overview | 96

Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment | 104

Configuring IGMP Snooping on Switches | 122

Example: Configuring IGMP Snooping on Switches | 126

Requirements | 126

Overview and Topology | 126

Configuration | 127

Example: Configuring IGMP Snooping on EX Series Switches | 130

Requirements | 130

Overview and Topology | 130

Configuration | 132

Verifying IGMP Snooping Operation | 133

Verifying IGMP Snooping on EX Series Switches | 134

Verifying IGMP Snooping Memberships | 135

Viewing IGMP Snooping Statistics | 136

Viewing IGMP Snooping Routing Information | 137

Changing the IGMP Snooping Group Timeout Value on Switches | 138

Monitoring IGMP Snooping | 139

Example: Configuring IGMP Snooping | 141

Understanding Multicast Snooping | 142

Understanding IGMP Snooping | 142

IGMP Snooping Interfaces and Forwarding | 144

IGMP Snooping and Proxies | 144

Multicast-Router Interfaces and IGMP Snooping Proxy Mode | 145

Host-Side Interfaces and IGMP Snooping Proxy Mode | 146

IGMP Snooping and Bridge Domains | 146

Configuring IGMP Snooping | 147

Configuring VLAN-Specific IGMP Snooping Parameters | 148

Example: Configuring IGMP Snooping | 149

Requirements | 149

Overview and Topology | 150

Configuration | 154

Verification | 158

Configuring IGMP Snooping Trace Operations | 158

Configuring IGMP or MLD Snooping Version | 161

Restrict Flooding of Unknown Multicast Traffic to the Multicast-router Interface in an L2 environment | 163

Example: Configuring IGMP Snooping on SRX Series Devices | 167

Requirements | 167

Overview and Topology | 168

Configuration | 169

Verifying IGMP Snooping Operation | 173

Configuring Point-to-Multipoint LSP with IGMP Snooping | 174

Configuring MLD Snooping | 177

Understanding MLD Snooping | 177

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure) | 188

Enabling or Disabling MLD Snooping on VLANs | 190

Configuring the MLD Version | 191

Enabling Immediate Leave | 192

Configuring an Interface as a Multicast-Router Interface | 193

Configuring Static Group Membership on an Interface | 194

Changing the Timer and Counter Values | 195

Configuring MLD Snooping on a Switch VLAN with ELS Support (CLI Procedure) | 197

Enabling or Disabling MLD Snooping on VLANs | 198

Configuring the MLD Version | 199

Enabling Immediate Leave | 200

Configuring an Interface as a Multicast-Router Interface | 200

Configuring Static Group Membership on an Interface | 201

Changing the Timer and Counter Values | 202

Example: Configuring MLD Snooping on EX Series Switches | 204

Requirements | 204

Overview and Topology | 205

Configuration | 207

Verifying MLD Snooping Configuration | 208

Example: Configuring MLD Snooping on SRX Series Devices | 209

Requirements | 210

Overview and Topology | 210

Configuration | 211

Verifying MLD Snooping Configuration | 215

Configuring MLD Snooping Tracing Operations on EX Series Switches (CLI Procedure) | 217

Configuring Tracing Operations | 218

Viewing, Stopping, and Restarting Tracing Operations | 219

Configuring MLD Snooping Tracing Operations on EX Series Switch VLANs (CLI Procedure) | 220

Configuring Tracing Operations | 221

Viewing, Stopping, and Restarting Tracing Operations | 222

Example: Configuring MLD Snooping on EX Series Switches | 223

Requirements | 223

Overview and Topology | 224

Configuration | 226

Verifying MLD Snooping Configuration | 227

Example: Configuring MLD Snooping on Switches with ELS Support | 228

Requirements | 229

Overview and Topology | 229

Configuration | 231

Verifying MLD Snooping Configuration | 232

Verifying MLD Snooping on EX Series Switches (CLI Procedure) | 234

Verifying MLD Snooping Memberships | 234

Verifying MLD Snooping VLANs | 235

Viewing MLD Snooping Statistics | 236

Viewing MLD Snooping Routing Information | 237

Verifying MLD Snooping on Switches | 239

Verifying MLD Snooping Memberships | 239

Verifying MLD Snooping Interfaces | 240

Viewing MLD Snooping Statistics | 242

Viewing MLD Snooping Routing Information | 243

Configuring Multicast VLAN Registration | 245

Understanding Multicast VLAN Registration | 245

Configuring Multicast VLAN Registration on EX Series Switches | 255

Configuring Multicast VLAN Registration on EX Series Switches with ELS | 256

Viewing MVLAN and MVR Receiver VLAN Information on EX Series Switches with ELS | 265

Configuring Multicast VLAN Registration on non-ELS EX Series Switches | 266

Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS | 267

Requirements | 268

Overview and Topology | 268

Configuration | 271

3

Configuring Protocol Independent Multicast

Understanding PIM | 275

PIM Overview | 275

PIM on Aggregated Interfaces | 279

Configuring PIM Basics | 280

Configuring Different PIM Modes | 280

Configuring Multiple Instances of PIM | 282

Changing the PIM Version | 282

Optimizing the Number of Multicast Flows on QFabric Systems | 283

Modifying the PIM Hello Interval | 283

Preserving Multicast Performance by Disabling Response to the ping Utility | 285

Configuring PIM Trace Options | 286

Configuring BFD for PIM | 289

Configuring BFD Authentication for PIM | 291

 Configuring BFD Authentication Parameters | 291

 Viewing Authentication Information for BFD Sessions | 293

Routing Content to Densely Clustered Receivers with PIM Dense Mode | 296

Understanding PIM Dense Mode | 296

Understanding PIM Sparse-Dense Mode | 298

Mixing PIM Sparse and Dense Modes | 299

Configuring PIM Dense Mode | 299

 Understanding PIM Dense Mode | 299

 Configuring PIM Dense Mode Properties | 302

Configuring PIM Sparse-Dense Mode | 304

 Understanding PIM Sparse-Dense Mode | 304

 Mixing PIM Sparse and Dense Modes | 304

 Configuring PIM Sparse-Dense Mode Properties | 305

Routing Content to Larger, Sparser Groups with PIM Sparse Mode | 307

Understanding PIM Sparse Mode | 307

Examples: Configuring PIM Sparse Mode | 311

 Understanding PIM Sparse Mode | 311

 Understanding Designated Routers | 314

 Tunnel Services PICs and Multicast | 315

 Enabling PIM Sparse Mode | 316

 Configuring PIM Join Load Balancing | 317

 Modifying the Join State Timeout | 321

 Example: Enabling Join Suppression | 322

 Requirements | 322

 Overview | 323

 Configuration | 325

 Verification | 328

 Example: Configuring PIM Sparse Mode over an IPsec VPN | 328

 Example: Configuring Multicast for Virtual Routers with IPv6 Interfaces | 335

 Requirements | 335

 Overview | 335

Configuration | 336

Verification | 341

Configuring Static RP | 342

Understanding Static RP | 342

Configuring Local PIM RPs | 343

Example: Configuring PIM Sparse Mode and RP Static IP Addresses | 345

Requirements | 346

Overview | 346

Configuration | 346

Verification | 348

Configuring the Static PIM RP Address on the Non-RP Routing Device | 350

Example: Configuring Anycast RP | 352

Understanding RP Mapping with Anycast RP | 353

Example: Configuring Multiple RPs in a Domain with Anycast RP | 353

Requirements | 354

Overview | 354

Configuration | 354

Verification | 357

Example: Configuring PIM Anycast With or Without MSDP | 358

Configuring a PIM Anycast RP Router Using Only PIM | 362

Configuring PIM Bootstrap Router | 365

Understanding the PIM Bootstrap Router | 365

Configuring PIM Bootstrap Properties for IPv4 | 366

Configuring PIM Bootstrap Properties for IPv4 or IPv6 | 367

Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain | 369

Example: Configuring PIM BSR Filters | 370

Understanding PIM Auto-RP | 371

Configuring All PIM Anycast Non-RP Routers | 371

Configuring a PIM Anycast RP Router with MSDP | 372

Configuring Embedded RP | 373

Understanding Embedded RP for IPv6 Multicast | 373

Configuring PIM Embedded RP for IPv6 | 375

Configuring PIM Filtering	376
Understanding Multicast Message Filters	377
Filtering MAC Addresses	378
Filtering RP and DR Register Messages	378
Filtering MSDP SA Messages	379
Configuring Interface-Level PIM Neighbor Policies	380
Filtering Outgoing PIM Join Messages	381
Example: Stopping Outgoing PIM Register Messages on a Designated Router	382
Requirements	382
Overview	383
Configuration	383
Verification	386
Filtering Incoming PIM Join Messages	387
Example: Rejecting Incoming PIM Register Messages on RP Routers	389
Requirements	389
Overview	390
Configuration	390
Verification	393
Configuring Register Message Filters on a PIM RP and DR	394
Examples: Configuring PIM RPT and SPT Cutover	397
Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees	398
Building an RPT Between the RP and Receivers	399
PIM Sparse Mode Source Registration	400
Multicast Shortest-Path Tree	403
SPT Cutover	404
SPT Cutover Control	409
Example: Configuring the PIM Assert Timeout	410
Requirements	410
Overview	410
Configuration	412
Example: Configuring the PIM SPT Threshold Policy	413
Requirements	414
Overview	414
Configuration	416
Verification	418

Disabling PIM | 419

- Disabling the PIM Protocol | 419
- Disabling PIM on an Interface | 420
- Disabling PIM for a Family | 421
- Disabling PIM for a Rendezvous Point | 422

Configuring Designated Routers | 424

Understanding Designated Routers | 424

Configuring a Designated Router for PIM | 425

- Configuring Interface Priority for PIM Designated Router Selection | 425
- Configuring PIM Designated Router Election on Point-to-Point Links | 427

Receiving Content Directly from the Source with SSM | 429

Understanding PIM Source-Specific Mode | 429

Example: Configuring Source-Specific Multicast | 434

- Understanding PIM Source-Specific Mode | 434
- Source-Specific Multicast Groups Overview | 438
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 439
 - Requirements | 439
 - Overview | 440
 - Configuration | 442
 - Verification | 444

Example: Configuring an SSM-Only Domain | 445

Example: Configuring PIM SSM on a Network | 445

Example: Configuring SSM Mapping | 448

Example: Configuring PIM SSM on a Network | 451

Example: Configuring an SSM-Only Domain | 454

Example: Configuring SSM Mapping | 455

Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 458

- Requirements | 459
- Overview | 459
- Configuration | 461
- Verification | 463

Example: Configuring SSM Maps for Different Groups to Different Sources | 464

Multiple SSM Maps and Groups for Interfaces | 464

Example: Configuring Multiple SSM Maps Per Interface | 464

Requirements | 464

Overview | 464

Configuration | 465

Verification | 468

Minimizing Routing State Information with Bidirectional PIM | 470

Example: Configuring Bidirectional PIM | 470

Understanding Bidirectional PIM | 470

Example: Configuring Bidirectional PIM | 478

Requirements | 478

Overview | 478

Configuration | 481

Verification | 488

Rapidly Detecting Communication Failures with PIM and the BFD Protocol | 497

Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol | 497

Understanding Bidirectional Forwarding Detection Authentication for PIM | 497

Configuring BFD for PIM | 500

Configuring BFD Authentication for PIM | 502

Configuring BFD Authentication Parameters | 502

Viewing Authentication Information for BFD Sessions | 504

Example: Configuring BFD Liveness Detection for PIM IPv6 | 506

Requirements | 507

Overview | 507

Configuration | 508

Verification | 513

Configuring PIM Options | 515

Example: Configuring Nonstop Active Routing for PIM | 515

Understanding Nonstop Active Routing for PIM | 515

Example: Configuring Nonstop Active Routing with PIM | 516

Requirements | 516

Overview | 517

Configuration | 518

Verification | 531

Configuring PIM Sparse Mode Graceful Restart | 532

Configuring PIM-to-IGMP and PIM-to-MLD Message Translation | 534

Understanding PIM-to-IGMP and PIM-to-MLD Message Translation | 534

Configuring PIM-to-IGMP Message Translation | 535

Configuring PIM-to-MLD Message Translation | 537

Verifying PIM Configurations | 539

Verifying the PIM Mode and Interface Configuration | 539

Verifying the PIM RP Configuration | 540

Verifying the RPF Routing Table Configuration | 541

Configuring Multicast Routing Protocols

Connecting Routing Domains Using MSDP | 544

Examples: Configuring MSDP | 544

Understanding MSDP | 544

Configuring MSDP | 546

Example: Configuring MSDP in a Routing Instance | 548

Requirements | 548

Overview | 548

Configuration | 551

Verification | 557

Configuring the Interface to Accept Traffic from a Remote Source | 557

Example: Configuring MSDP with Active Source Limits and Mesh Groups | 559

Requirements | 559

Overview | 559

Configuration | 563

Verification | 566

Tracing MSDP Protocol Traffic | 566

Disabling MSDP | 569

Example: Configuring MSDP | 570

Configuring Multiple Instances of MSDP | 571

Handling Session Announcements with SAP and SDP | 573

Configuring the Session Announcement Protocol | 573

Understanding SAP and SDP | 573

Configuring the Session Announcement Protocol | 574

Verifying SAP and SDP Addresses and Ports | 575

Facilitating Multicast Delivery Across Unicast-Only Networks with AMT | 577

Example: Configuring Automatic IP Multicast Without Explicit Tunnels | 577

Understanding AMT | 577

AMT Applications | 579

AMT Operation | 580

Configuring the AMT Protocol | 581

Configuring Default IGMP Parameters for AMT Interfaces | 584

Example: Configuring the AMT Protocol | 587

Requirements | 588

Overview | 588

Configuration | 589

Verification | 592

Routing Content to Densely Clustered Receivers with DVMRP | 594

Examples: Configuring DVMRP | 594

Understanding DVMRP | 594

Configuring DVMRP | 595

Example: Configuring DVMRP | 596

Requirements | 596

Overview | 596

Configuration | 598

Verification | 600

Example: Configuring DVMRP to Announce Unicast Routes | 600

Requirements | 601

Overview | 601

Configuration | 602

Verification | 605

Tracing DVMRP Protocol Traffic | 606

5

Configuring Multicast VPNs

Configuring Draft-Rosen Multicast VPNs | 610

Draft-Rosen Multicast VPNs Overview | 610

Example: Configuring Any-Source Draft-Rosen 6 Multicast VPNs | **611**

Understanding Any-Source Multicast | **611**

Example: Configuring Any-Source Multicast for Draft-Rosen VPNs | **612**

Requirements | **612**

Overview | **613**

Configuration | **616**

Verification | **625**

Load Balancing Multicast Tunnel Interfaces Among Available PICs | **625**

Example: Configuring a Specific Tunnel for IPv4 Multicast VPN Traffic (Using Draft-Rosen MVPNs) | **630**

Requirements | **630**

Overview | **631**

PE Router Configuration | **632**

CE Device Configuration | **641**

Verification | **644**

Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs | **648**

Understanding Source-Specific Multicast VPNs | **648**

Draft-Rosen 7 Multicast VPN Control Plane | **649**

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs | **649**

Requirements | **649**

Overview | **650**

Configuration | **652**

Verification | **660**

Understanding Data MDTs | **661**

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode | **663**

Requirements | **663**

Overview | **663**

Configuration | **666**

Verification | **668**

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode | **669**

Requirements | **669**

Overview | **670**

Configuration | **676**

Verification | **681**

Examples: Configuring Data MDTs | **682**

Understanding Data MDTs | **683**

Data MDT Characteristics | **684**

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode | **685**

Requirements | **685**

Overview | **686**

Configuration | **692**

Verification | **697**

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode | **698**

Requirements | **699**

Overview | **699**

Configuration | **702**

Verification | **703**

Example: Enabling Dynamic Reuse of Data MDT Group Addresses | **704**

Requirements | **704**

Overview | **705**

Configuration | **706**

Verification | **714**

Configuring Next-Generation Multicast VPNs | 715

Understanding Next-Generation MVPN Network Topology | **716**

Understanding Next-Generation MVPN Concepts and Terminology | **718**

Understanding Next-Generation MVPN Control Plane | **720**

Next-Generation MVPN Data Plane Overview | **730**

Enabling Next-Generation MVPN Services | **734**

Generating Next-Generation MVPN VRF Import and Export Policies Overview | **737**

Multiprotocol BGP MVPNs Overview | **741**

Comparison of Draft Rosen Multicast VPNs and Next-Generation Multiprotocol BGP Multicast VPNs | **741**

MBGP Multicast VPN Sites | **742**

Multicast VPN Standards | **743**

PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs | 743

MBGP-Based Multicast VPN Trees | 744

Configuring Multiprotocol BGP Multicast VPNs | 749

Understanding Multiprotocol BGP-Based Multicast VPNs: Next-Generation | 750

Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs | 751

Requirements | 751

Overview | 753

Configuration | 755

Verification | 757

Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs | 758

Requirements | 758

Overview | 759

Configuration | 761

Verification | 766

Example: Configuring MBGP Multicast VPNs | 776

Requirements | 776

Overview and Topology | 776

Configuration | 777

Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN | 800

Requirements | 800

Overview | 800

Configuration | 802

Verification | 811

Example: Allowing MBGP MVPN Remote Sources | 812

Requirements | 812

Overview | 812

Configuration | 814

Verification | 818

Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family | 819

Requirements | 819

Overview | 819

Configuration | 820

Verification | 832

Example: Configuring MBGP Multicast VPN Topology Variations | 834

Requirements | **834**

Overview and Topology | **834**

Configuring Full Mesh MBGP MVPNs | **837**

Configuring Sender-Only and Receiver-Only Sites Using PIM ASM Provider Tunnels | **840**

Configuring Sender-Only, Receiver-Only, and Sender-Receiver MVPN Sites | **843**

Configuring Hub-and-Spoke MVPNs | **847**

Configuring Nonstop Active Routing for BGP Multicast VPN | **850**

BGP-MVPN Inter-AS Option B Overview | **854**

ACX Support for BGP MVPN | **856**

Example: Configuring MBGP MVPN Extranets | **858**

Understanding MBGP Multicast VPN Extranets | **858**

MBGP Multicast VPN Extranets Configuration Guidelines | **859**

Example: Configuring MBGP Multicast VPN Extranets | **860**

Requirements | **860**

Overview and Topology | **861**

Configuration | **862**

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs | **913**

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs | **914**

Requirements | **914**

Overview | **914**

Configuration | **915**

Verification | **925**

Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels | **928**

Multiple Active and Backup Paths in RPF List | **932**

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels | **935**

Requirements | **936**

Overview | **936**

Set Commands for All Devices in the Topology | **938**

Configuring Device PE2 | **943**

Verification | **952**

Example: Configuring Sender-Based RPF in a BGP MVPN with MLDP Point-to-Multipoint Provider Tunnels | **972**

Requirements | **972**

Overview | **973**

Set Commands for All Devices in the Topology | **974**

Configuring Device PE2 | **980**

Verification | **988**

Configuring MBGP MVPN Wildcards | **1008**

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN | **1008**

Configuring a Selective Provider Tunnel Using Wildcards | **1014**

Example: Configuring Selective Provider Tunnels Using Wildcards | **1015**

Distributing C-Multicast Routes Overview | **1017**

Exchanging C-Multicast Routes | **1023**

Generating Source AS and Route Target Import Communities Overview | **1032**

Originating Type 1 Intra-AS Autodiscovery Routes Overview | **1033**

Signaling Provider Tunnels and Data Plane Setup | **1037**

Anti-spoofing support for MPLS labels in BGP/MPLS IP VPNs (Inter-AS Option B) | **1055**

BGP-MVPN SD-WAN Overlay | **1056**

Configuring PIM Join Load Balancing | 1059

Use Case for PIM Join Load Balancing | **1059**

Configuring PIM Join Load Balancing | **1060**

PIM Join Load Balancing on Multipath MVPN Routes Overview | **1064**

Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN | **1068**

Requirements | **1069**

Overview and Topology | **1069**

Configuration | **1073**

Verification | **1078**

Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN | **1080**

Requirements | **1080**

Overview and Topology | **1081**

Configuration | 1084

Verification | 1090

Example: Configuring PIM Make-Before-Break Join Load Balancing | 1092

Understanding the PIM Automatic Make-Before-Break Join Load-Balancing Feature | 1092

Example: Configuring PIM Make-Before-Break Join Load Balancing | 1093

Requirements | 1093

Overview | 1094

Configuration | 1095

Verification | 1101

Example: Configuring PIM State Limits | 1106

Controlling PIM Resources for Multicast VPNs Overview | 1106

Example: Configuring PIM State Limits | 1109

Requirements | 1110

Overview | 1110

Configuration | 1111

Verification | 1122

6

General Multicast Options

Bit Index Explicit Replication (BIER) | 1125

Overview | 1125

IS-IS extension for BIER | 1129

Configuring BIER MVPN | 1132

BIER Forwarding | 1136

Prevent Routing Loops with Reverse Path Forwarding | 1149

Examples: Configuring Reverse Path Forwarding | 1149

Understanding Multicast Reverse Path Forwarding | 1149

Multicast RPF Configuration Guidelines | 1151

Example: Configuring a Dedicated PIM RPF Routing Table | 1152

Requirements | 1152

Overview | 1153

Configuration | 1154

Example: Configuring a PIM RPF Routing Table | 1157

Requirements | 1158

Overview | **1158**

Configuration | **1159**

Verification | **1161**

Example: Configuring RPF Policies | **1164**

Requirements | **1164**

Overview | **1164**

Configuration | **1165**

Verification | **1167**

Example: Configuring PIM RPF Selection | **1167**

Requirements | **1168**

Overview | **1168**

Configuration | **1170**

Verification | **1172**

Use Multicast-Only Fast Reroute (MoFRR) to Minimize Packet Loss During Link Failures | 1173

Understanding Multicast-Only Fast Reroute | **1173**

Configuring Multicast-Only Fast Reroute | **1182**

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain | **1185**

Requirements | **1186**

Overview | **1186**

CLI Quick Configuration | **1188**

Step-by-Step Configuration | **1190**

Verification | **1194**

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches | **1197**

Requirements | **1197**

Overview | **1198**

CLI Quick Configuration | **1199**

Step-by-Step Configuration | **1201**

Verification | **1205**

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | **1208**

Requirements | **1209**

Overview | **1209**

CLI Quick Configuration | **1210**

Configuration | **1219**

| Verification | 1226

Enable Multicast Between Layer 2 and Layer 3 Devices Using Snooping | 1232

Multicast Snooping on MX Series Routers | 1232

Example: Configuring Multicast Snooping | 1233

Understanding Multicast Snooping | 1233

Understanding Multicast Snooping and VPLS Root Protection | 1234

Configuring Multicast Snooping | 1235

Example: Configuring Multicast Snooping | 1236

Requirements | 1236

Overview and Topology | 1237

Configuration | 1239

Verification | 1242

Enabling Bulk Updates for Multicast Snooping | 1243

Enabling Multicast Snooping for Multichassis Link Aggregation Group Interfaces | 1244

Example: Configuring Multicast Snooping for a Bridge Domain | 1245

Configuring Multicast Snooping to Ignore Spanning Tree Topology Change Messages | 1246

Configuring Graceful Restart for Multicast Snooping | 1248

Multicast Snooping for VPLS | 1250

Understanding PIM Snooping for VPLS | 1250

Example: Configuring PIM Snooping for VPLS | 1252

Requirements | 1252

Overview | 1252

Configuration | 1253

Verification | 1263

IGMP and MLD Snooping for VPLS | 1267

Configure Multicast Routing Options | 1273

Examples: Configuring Administrative Scoping | 1273

Understanding Multicast Administrative Scoping | 1273

Example: Creating a Named Scope for Multicast Scoping | 1275

Requirements | 1275

Overview | 1276

Configuration | 1276

Verification | **1279**

Example: Using a Scope Policy for Multicast Scoping | **1279**

Requirements | **1279**

Overview | **1279**

Configuration | **1280**

Verification | **1283**

Example: Configuring Externally Facing PIM Border Routers | **1283**

Examples: Configuring Bandwidth Management | **1284**

Understanding Bandwidth Management for Multicast | **1284**

Bandwidth Management and PIM Graceful Restart | **1285**

Bandwidth Management and Source Redundancy | **1285**

Logical Systems and Bandwidth Oversubscription | **1285**

Example: Defining Interface Bandwidth Maximums | **1287**

Requirements | **1287**

Overview | **1287**

Configuration | **1289**

Verification | **1291**

Example: Configuring Multicast with Subscriber VLANs | **1291**

Requirements | **1291**

Overview and Topology | **1292**

Configuration | **1296**

Verification | **1307**

Configuring Multicast Routing over IP Demux Interfaces | **1308**

Classifying Packets by Egress Interface | **1309**

Recycle Bandwidth Management | **1312**

Example: Configuring Recycle Bandwidth | **1314**

Overview | **1314**

Configure Recycle Interface Bandwidth | **1315**

Verification | **1315**

Examples: Configuring the Multicast Forwarding Cache | **1317**

Understanding the Multicast Forwarding Cache | **1317**

Example: Configuring the Multicast Forwarding Cache | **1317**

Requirements | **1318**

Overview | **1318**

Configuration | **1319**

Verification | 1321

Example: Configuring a Multicast Flow Map | 1321

Requirements | 1322

Overview | 1322

Configuration | 1323

Verification | 1326

Example: Configuring Ingress PE Redundancy | 1327

Understanding Ingress PE Redundancy | 1327

Example: Configuring Ingress PE Redundancy | 1327

Requirements | 1328

Overview | 1328

Configuration | 1330

Verification | 1334

Controller-Based BGP Multicast Signaling | 1335

Controller-Based BGP Multicast Signaling | 1335

7

Troubleshooting

Knowledge Base | 1340

8

Configuration Statements and Operational Commands

Junos CLI Reference Overview | 1342

About This Guide

Multicast allows an IP network to support more than just the unicast model of data delivery that prevailed in the early stages of the Internet. Multicast provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many.

In a multicast network, the key component is the routing device, which is able to replicate packets and is therefore multicast-capable. The routing devices in the IP multicast network, which has exactly the same topology as the unicast network it is based on, use a multicast routing protocol to build a distribution tree that connects receivers (preferred to the multimedia implications of listeners, but listeners is also used) to sources. In multicast terminology, the distribution tree is rooted at the source (the root of the distribution tree is the source). The interface on the routing device leading toward the source is the upstream interface, although the less precise terms incoming or inbound interface are used as well. To keep bandwidth use to a minimum, it is best for only one upstream interface on the routing device to receive multicast packets. The interface on the routing device leading toward the receivers is the downstream interface, although the less precise terms outgoing or outbound interface are used as well. There can be 0 to N-1 downstream interfaces on a routing device, where N is the number of logical interfaces on the routing device.

RELATED DOCUMENTATION

[vDay One: Introduction to BGP Multicast VPNs](#)

[This Week: Deploying BGP Multicast VPNs](#)

1

PART

Overview

Understanding Multicast | 2

CHAPTER 1

Understanding Multicast

IN THIS CHAPTER

- [Multicast Overview | 2](#)
- [Understanding Layer 3 Multicast Functionality on the SRX5K-MPC | 16](#)
- [Multicast Configuration Overview | 17](#)
- [IPv6 Multicast Flow | 19](#)
- [Supported IP Multicast Protocol Standards | 21](#)
- [Ingress and Egress Multicast Replication using Recycle Replication Model | 23](#)

Multicast Overview

IN THIS SECTION

- [Comparing Multicast to Unicast | 3](#)
- [IP Multicast Uses | 5](#)
- [IP Multicast Terminology | 6](#)
- [Reverse-Path Forwarding for Loop Prevention | 7](#)
- [Shortest-Path Tree for Loop Prevention | 7](#)
- [Administrative Scoping for Loop Prevention | 7](#)
- [Multicast Leaf and Branch Terminology | 7](#)
- [IP Multicast Addressing | 8](#)
- [Multicast Addresses | 9](#)
- [Layer 2 Frames and IPv4 Multicast Addresses | 9](#)
- [Multicast Interface Lists | 12](#)
- [Multicast Routing Protocols | 13](#)
- [T Series Router Multicast Performance | 16](#)

IP has three fundamental types of addresses: unicast, broadcast, and multicast. A *unicast address* is used to send a packet to a single destination. A *broadcast address* is used to send a datagram to an entire subnetwork. A *multicast address* is used to send a datagram to a set of hosts that can be on different subnetworks and that are configured as members of a multicast group.

A multicast datagram is delivered to destination group members with the same best-effort reliability as a standard unicast IP datagram. This means that multicast datagrams are not guaranteed to reach all members of a group or to arrive in the same order in which they were transmitted. The only difference between a multicast IP packet and a unicast IP packet is the presence of a group address in the IP header destination address field. Multicast addresses use the Class D address format.



NOTE: On all SRX Series Firewalls, reordering is not supported for multicast fragments. Reordering of unicast fragments is supported.

Individual hosts can join or leave a multicast group at any time. There are no restrictions on the physical location or the number of members in a multicast group. A host can be a member of more than one multicast group at any time. A host does not have to belong to a group to send packets to members of a group.

Routers use a group membership protocol to learn about the presence of group members on directly attached subnetworks. When a host joins a multicast group, it transmits a group membership protocol message for the group or groups that it wants to receive and sets its IP process and network interface card to receive frames addressed to the multicast group.

Comparing Multicast to Unicast

The Junos[®] operating system (Junos OS) routing protocol process supports a wide variety of routing protocols. These routing protocols carry network information among routing devices not only for *unicast* traffic streams sent between one pair of clients and servers, but also for *multicast* traffic streams containing video, audio, or both, between a single server source and many client receivers. The routing protocols used for multicast differ in many key ways from unicast routing protocols.

Information is delivered over a network by three basic methods: unicast, broadcast, and multicast.

The differences among unicast, broadcast, and multicast can be summarized as follows:

- Unicast: One-to-one, from one source to one destination.
- Broadcast: One-to-all, from one source to all possible destinations.
- Multicast: One-to-many, from one source to multiple destinations expressing an interest in receiving the traffic.



NOTE: This list does not include a special category for many-to-many applications, such as online gaming or videoconferencing, where there are many sources for the same receiver and where receivers often double as sources. Many-to-many is a service model that repeatedly employs one-to-many multicast and therefore requires no unique protocol. The original multicast specification, RFC 1112, supports both the any-source multicast (ASM) many-to-many model and the source-specific multicast (SSM) one-to-many model.

With unicast traffic, many streams of IP packets that travel across networks flow from a single source, such as a website server, to a single destination such as a client PC. Unicast traffic is still the most common form of information transfer on networks.

Broadcast traffic flows from a single source to all possible destinations reachable on the network, which is usually a LAN. Broadcasting is the easiest way to make sure traffic reaches its destinations.

Television networks use broadcasting to distribute video and audio. Even if the television network is a cable television (CATV) system, the source signal reaches all possible destinations, which is the main reason that some channels' content is scrambled. Broadcasting is not feasible on the Internet because of the enormous amount of unnecessary information that would constantly arrive at each end user's device, the complexities and impact of scrambling, and related privacy issues.

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

On an IP network, because destinations (clients) do not often communicate directly with sources (servers), the routing devices between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. Multicast routing devices replicate packets received on one input interface and send the copies out on multiple output interfaces.

In IP multicast, the source and destination are almost always hosts and not routing devices. Multicast routing devices distribute the multicast traffic across the network from source to destinations. The multicast routing device must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities, but some router architectures cannot send multiple copies of packets and so do not support multicasting directly.

IP Multicast Uses

Multicast allows an IP network to support more than just the unicast model of data delivery that prevailed in the early stages of the Internet. Multicast, originally defined as a host extension in RFC 1112 in 1989, provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many.

Unicast traffic is not strictly limited to data applications. Telephone conversations, wireless or not, contain digital audio samples and might contain digital photographs or even video and still flow from a single source to a single destination. In the same way, multicast traffic is not strictly limited to multimedia applications. In some data applications, the flow of traffic is from a single source to many destinations that require the packets, as in a news or stock ticker service delivered to many PCs. For this reason, the term *receiver* is preferred to *listener* for multicast destinations, although both terms are common.

Network applications that can function with unicast but are better suited for multicast include collaborative groupware, teleconferencing, periodic or “push” data delivery (stock quotes, sports scores, magazines, newspapers, and advertisements), server or website replication, and distributed interactive simulation (DIS) such as war simulations or virtual reality. Any IP network concerned with reducing network resource overhead for one-to-many or many-to-many data or multimedia applications with multiple receivers benefits from multicast.

If unicast were employed by radio or news ticker services, each radio or PC would have to have a separate traffic session for each listener or viewer at a PC (this is actually the method for some Web-based services). The processing load and bandwidth consumed by the server would increase linearly as more people “tune in” to the server. This is extremely inefficient when dealing with the global scale of the Internet. Unicast places the burden of packet duplication on the server and consumes more and more backbone bandwidth as the number of users grows.

If broadcast were employed instead, the source could generate a single IP packet stream using a broadcast destination address. Although broadcast eliminates the server packet duplication issue, this is not a good solution for IP because IP broadcasts can be sent only to a single subnetwork, and IP routing devices normally isolate IP subnetworks on separate interfaces. Even if an IP packet stream could be addressed to literally go everywhere, and there were no need to “tune” to any source at all, broadcast would be extremely inefficient because of the bandwidth strain and need for uninterested hosts to discard large numbers of packets. Broadcast places the burden of packet rejection on each host and consumes the maximum amount of backbone bandwidth.

For radio station or news ticker traffic, multicast provides the most efficient and effective outcome, with none of the drawbacks and all of the advantages of the other methods. A single source of multicast packets finds its way to every *interested* receiver. As with broadcast, the transmitting host generates only a single stream of IP packets, so the load remains constant whether there is one receiver or one million. The network routing devices replicate the packets and deliver the packets to the proper receivers, but only the replication role is a new one for routing devices. The links leading to subnets

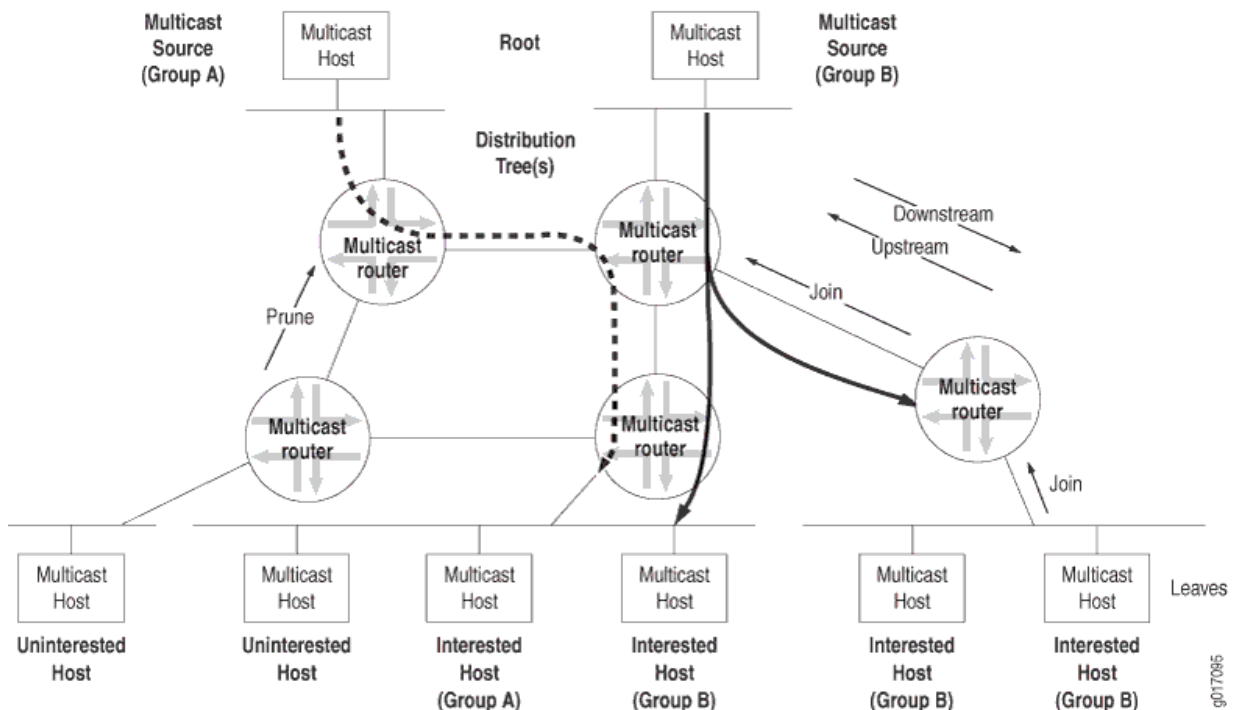
consisting of entirely uninterested receivers carry no multicast traffic. Multicast minimizes the burden placed on sender, network, and receiver.

IP Multicast Terminology

Multicast has its own particular set of terms and acronyms that apply to IP multicast routing devices and networks. Figure 1 on page 6 depicts some of the terms commonly used in an IP multicast network.

In a multicast network, the key component is the *routing device*, which is able to replicate packets and is therefore multicast-capable. The routing devices in the IP multicast network, which has exactly the same topology as the unicast network it is based on, use a *multicast routing protocol* to build a *distribution tree* that connects receivers (preferred to the multimedia implications of listeners, but listeners is also used) to *sources*. In multicast terminology, the distribution tree is *rooted at the source* (the root of the distribution tree is the source). The interface on the routing device leading toward the source is the *upstream* interface, although the less precise terms *incoming* or *inbound* interface are used as well. To keep bandwidth use to a minimum, it is best for only one upstream interface on the routing device to receive multicast packets. The interface on the routing device leading toward the receivers is the *downstream* interface, although the less precise terms *outgoing* or *outbound* interface are used as well. There can be 0 to $N-1$ downstream interfaces on a routing device, where N is the number of logical interfaces on the routing device. To prevent looping, the upstream interface must never receive copies of downstream multicast packets.

Figure 1: Multicast Terminology in an IP Network



Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols.

Reverse-Path Forwarding for Loop Prevention

The routing device's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In RPF, every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface.

When a router receives a multicast packet on an interface, the router verifies the *source* address. The router then checks to see if the same address could be used as the *destination* address to get back to the source via a unicast route. If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check.

Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. Routing devices can build and maintain separate tables for RPF purposes.

Shortest-Path Tree for Loop Prevention

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast routing device operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

Administrative Scoping for Loop Prevention

Scoping limits the routing devices and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. Routing devices at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

Multicast Leaf and Branch Terminology

Each subnetwork with hosts on the routing device that has at least one interested receiver is a *leaf* on the distribution tree. Routing devices can have multiple leaves on different interfaces and must send a copy of the IP multicast packet out on each interface with a leaf. When a new leaf subnetwork is added to the tree (that is, the interface to the host subnetwork previously received no copies of the multicast packets), a new *branch* is built, the leaf is joined to the tree, and replicated packets are sent out on the

interface. The number of leaves on a particular interface does not affect the routing device. The action is the same for one leaf or a hundred.



NOTE: On Juniper Networks security devices, if the maximum number of leaves on a multicast distribution tree is exceeded, multicast sessions are created up to the maximum number of leaves, and any multicast sessions that exceed the maximum number of leaves are ignored. The maximum number of leaves on a multicast distribution tree is device specific.

When a branch contains no leaves because there are no interested hosts on the routing device interface leading to that IP subnetwork, the branch is *pruned* from the distribution tree, and no multicast packets are sent out that interface. Packets are replicated and sent out multiple interfaces only where the distribution tree branches at a routing device, and no link ever carries a duplicate flow of packets.

Collections of hosts all receiving the same stream of IP packets, usually from the same multicast source, are called *groups*. In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast address, or *group address*. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

IP Multicast Addressing

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Class D addresses are commonly referred to as *multicast addresses* because the entire classful address concept is obsolete. Multicast addresses can never appear as the source address in an IP packet and can only be the destination of a packet.

Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices. Blocks of multicast addresses can still be described in terms of prefix length in traditional notation, but only for convenience. For example, the multicast address range from 232.0.0.0 through 232.255.255.255 can be written as 232.0.0.0/8 or 232/8.

Internet service providers (ISPs) do not typically allocate multicast addresses to their customers because multicast addresses relate to content, not to physical devices. Receivers are not assigned their own multicast addresses, but need to know the multicast address of the content. Sources need to be assigned multicast addresses only to produce the content, not to identify their place in the network. Every source and receiver still needs an ordinary, unicast IP address.

Multicast addressing most often references the receivers, and the source of multicast content is usually not even a member of the multicast group for which it produces content. If the source needs to monitor the packets it produces, monitoring can be done locally, and there is no need to make the packets traverse the network.

Many applications have been assigned a range of multicast addresses for their own use. These applications assign multicast addresses to sessions created by that application. You do not usually need to statically assign a multicast address, but you can do so.

Multicast Addresses

Multicast host group addresses are defined to be the IP addresses whose high-order four bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. (These addresses also are referred to as Class D addresses.)

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP multicast groups. The base address 224.0.0.0 is reserved and cannot be assigned to any group. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms.

The range from 239.0.0.0 through 239.255.255.255 is reserved for administratively scoped addresses. Because packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and because administratively scoped multicast addresses are locally assigned, these addresses do not need to be unique across administrative boundaries.

Layer 2 Frames and IPv4 Multicast Addresses

Multicasting on a LAN is a good place to start an investigation of multicasting at Layer 2. At Layer 2, multicast deals with media access control (MAC) frames and addresses instead of IPv4 or IPv6 packets and addresses. Consider a single LAN, without routing devices, with a multicast source sending to a certain group. The rest of the hosts are receivers interested in the multicast group's content. So the multicast source host generates packets with its unicast IP address as the source, and the multicast group address as the destination.

Which MAC addresses are used on the frame containing this packet? The packet source address—the unicast IP address of the host originating the multicast content—translates easily and directly to the MAC address of the source. But what about the packet's destination address? This is the IP multicast group address. Which destination MAC address for the frame corresponds to the packet's multicast group address?

One option is for LANs simply to use the LAN broadcast MAC address, which guarantees that the frame is processed by every station on the LAN. However, this procedure defeats the whole purpose of multicast, which is to limit the circulation of packets and frames to interested hosts. Also, hosts might have access to many multicast groups, which multiplies the amount of traffic to noninterested destinations. Broadcasting frames at the LAN level to support multicast groups makes no sense.

However, there is an easy way to effectively use Layer 2 frames for multicast purposes. The MAC address has a bit that is set to 0 for unicast (the LAN term is *individual address*) and set to 1 to indicate that this is a multicast address. Some of these addresses are reserved for multicast groups of specific

vendors or MAC-level protocols. Internet multicast applications use the range 0x01-00-5E-00-00-00 to 0x01-00-5E-FF-FF-FF. Multicast receivers (hosts running TCP/IP) listen for frames with one of these addresses when the application joins a multicast group. The host stops listening when the application terminates or the host leaves the group at the packet layer (Layer 3).

This means that 3 bytes, or 24 bits, are available to map IPv4 multicast addresses at Layer 3 to MAC multicast addresses at Layer 2. However, all IPv4 addresses, including multicast addresses, are 32 bits long, leaving 8 IP address bits left over. Which method of mapping IPv4 multicast addresses to MAC multicast addresses minimizes the chance of “collisions” (that is, two different IP multicast groups at the packet layer mapping to the same MAC multicast address at the frame layer)?

First, it is important to realize that all IPv4 multicast addresses begin with the same 4 bits (**1110**), so there are really only 4 bits of concern, not 8. A LAN must not drop the last bits of the IPv4 address because these are almost guaranteed to be host bits, depending on the subnet mask. But the high-order bits, the leftmost address bits, are almost always network bits, and there is only one LAN (for now).

One other bit of the remaining 24 MAC address bits is reserved (an initial **0** indicates an Internet multicast address), so the 5 bits following the initial 1110 in the IPv4 address are dropped. The 23 remaining bits are mapped, one for one, into the last 23 bits of the MAC address. An example of this process is shown in [Figure 2 on page 11](#).

Figure 2: Converting MAC Addresses to Multicast Addresses

1	IPv4 header multicast destination address	232.	224.	202.	181
	Written in hexadecimal	E8	E0	CA	B5
	Written in binary	1110 1000 1	1110 0000	1100 1010	1011 0101
2	Ignore the first 9 bits and copy the remaining 23 bits	X	X110 0000	1100 1010	1011 0101
3	First bit X = 0 for Internet; X = 1 for other	0	0110 0000	1100 1010	1011 0101
4	Written in hexadecimal		60	CA	B5
5	MAC address in hexadecimal	01:00:5E:E0:CA:B5			
6	Drop last 24 bits	01:00:5E:			
7	Copy the multicast bits	01:00:5E:60:CA:B5			
8	MAC frame destination address 01:00:5E:60:CA:B5 corresponds to multicast IPv4 address 232.224.202.181				

659910108

Note that this process means that there are 32 (2^5) IPv4 multicast addresses that could map to the same MAC multicast addresses. For example, multicast IPv4 addresses **224.8.7.6** and **229.136.7.6** translate to the same MAC address (0x01-00-5E-08-07-06). This is a real concern, and because the host could be interested in frames sent to both of the those multicast groups, the IP software must reject one or the other.



NOTE: This “collision” problem does not exist in IPv6 because of the way IPv6 handles multicast groups, but it is always a concern in IPv4. The procedure for placing IPv6 multicast packets inside multicast frames is nearly identical to that for IPv4, except for the MAC destination address 0x3333 prefix (and the lack of “collisions”).

Once the MAC address for the multicast group is determined, the host's operating system essentially orders the LAN interface card to join or leave the multicast group. Once joined to a multicast group, the host accepts frames sent to the multicast address as well as the host's unicast address and ignores other multicast group's frames. It is possible for a host to join and receive multicast content from more than one group at the same time, of course.

Multicast Interface Lists

To avoid multicast routing loops, every multicast routing device must always be aware of the interface that leads to the source of that multicast group content by the shortest path. This is the upstream (incoming) interface, and packets are never to be forwarded back toward a multicast source. All other interfaces are potential downstream (outgoing) interfaces, depending on the number of branches on the distribution tree.

Routing devices closely monitor the status of the incoming and outgoing interfaces, a process that determines the *multicast forwarding state*. A routing device with a multicast forwarding state for a particular multicast group is essentially “turned on” for that group's content. Interfaces on the routing device's outgoing interface list send copies of the group's packets received on the incoming interface list for that group. The incoming and outgoing interface lists might be different for different multicast groups.

The multicast forwarding state in a routing device is usually written in either (S,G) or (*,G) notation. These are pronounced “ess comma gee” and “star comma gee,” respectively. In (S,G), the S refers to the unicast IP address of the source for the multicast traffic, and the G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.

The asterisk (*) in the (*,G) notation is a wildcard indicating that the state applies to any multicast application source sending to group G. So, if two sources are originating exactly the same content for multicast group 224.1.1.2, a routing device could use (*,224.1.1.2) to represent the state of a routing device forwarding traffic from both sources to the group.

Multicast Routing Protocols

Multicast routing protocols enable a collection of multicast routing devices to build (join) distribution trees when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group, prune branches, locate sources and groups, and prevent routing loops.

There are several multicast routing protocols:

- **Distance Vector Multicast Routing Protocol (DVMRP)**—The first of the multicast routing protocols and hampered by a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G), and builds its own multicast routing tables for RPF checks.
- **Multicast OSPF (MOSPF)**—Extends OSPF for multicast use, but only for dense mode. However, MOSPF has an explicit join message, so routing devices do not have to flood their entire domain with multicast traffic from every source. MOSPF uses source-based distribution trees in the form (S,G).
- ***Bidirectional PIM mode***—A variation of PIM. Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees as in PIM-SM and is therefore optimized for routing state size instead of path length. This means that the end-to-end latency might be longer compared to PIM sparse mode. Bidirectional PIM routes are always wildcard-source (*,G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (*,G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM (*,G) routes forward traffic from all sources and the RP. Bidirectional PIM routing devices must have the ability to accept traffic on many potential incoming interfaces. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Bidirectional PIM is recommended in deployments with many dispersed sources and many dispersed receivers.
- **PIM *dense mode***—In this mode of PIM, the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is *flooded* with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). This is the *dense mode* of multicast operation. LANs are appropriate networks for dense-mode operation. Some multicast routing protocols, especially older ones, support only dense-mode operation, which makes them inappropriate for use on the Internet. In contrast to DVMRP and MOSPF, PIM dense mode allows a routing device to use any unicast routing protocol and performs RPF checks using the unicast routing table. PIM dense mode has an implicit join message, so routing devices use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are. PIM dense mode uses source-based distribution trees in the form (S,G), as do all dense-mode protocols. PIM also supports sparse-dense mode, with mixed sparse and dense groups, but there is no special

notation for that operational mode. If *sparse-dense mode* is supported, the multicast routing protocol allows some multicast groups to be sparse and other groups to be dense.

- PIM *sparse mode*—In this mode of PIM, the assumption is that very few of the possible receivers want packets from each source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) an interest in the traffic. This multicast protocol allows a routing device to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode has an *explicit* join message, so routing devices determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to the rendezvous point (RP). PIM sparse mode uses an RP routing device as the initial source of multicast group traffic and therefore builds distribution trees in the form (*,G), as do all sparse-mode protocols. PIM sparse mode migrates to an (S,G) source-based tree if that path is shorter than through the RP for a particular multicast group's traffic. WANs are appropriate networks for sparse-mode operation, and indeed a common multicast guideline is not to run dense mode on a WAN under any circumstances.
- Core Based Trees (CBT)—Shares all of the characteristics of PIM sparse mode (sparse mode, explicit join, and shared (*,G) trees), but is said to be more efficient at finding sources than PIM sparse mode. CBT is rarely encountered outside academic discussions. There are no large-scale deployments of CBT, commercial or otherwise.
- PIM source-specific multicast (SSM)—Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of an RP. Used with IGMPv3 to create a shortest-path tree between receiver and source.
- IGMPv1—The original protocol defined in RFC 1112, *Host Extensions for IP Multicasting*. IGMPv1 sends an explicit join message to the routing device, but uses a timeout to determine when hosts leave a group. Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routing devices.
- IGMPv2—Defined in RFC 2236, *Internet Group Management Protocol, Version 2*. Among other features, IGMPv2 adds an explicit leave message to the join message.
- IGMPv3—Defined in RFC 3376, *Internet Group Management Protocol, Version 3*. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or source-specific multicast (SSM). Used with PIM SSM to create a shortest-path tree between receiver and source.
- Bootstrap Router (BSR) and Auto-Rendezvous Point (RP)—Allow sparse-mode routing protocols to find RPs within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.
- Multicast Source Discovery Protocol (MSDP)—Allows groups located in one multicast routing domain to find RPs in other routing domains. MSDP is not used on an RP if all receivers and sources are

located in the same routing domain. Typically runs on the same routing device as PIM sparse mode RP. Not appropriate if all receivers and sources are located in the same routing domain.

- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.
- Pragmatic General Multicast (PGM)—Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.

The differences among the multicast routing protocols are summarized in [Table 1 on page 15](#).

Table 1: Multicast Routing Protocols Compared

Multicast Routing Protocol	Dense Mode	Sparse Mode	Implicit Join	Explicit Join	(S,G) SBT	(* ,G) Shared Tree
DVMRP	Yes	No	Yes	No	Yes	No
MOSPF	Yes	No	No	Yes	Yes	No
PIM dense mode	Yes	No	Yes	No	Yes	No
PIM sparse mode	No	Yes	No	Yes	Yes, maybe	Yes, initially
Bidirectional PIM	No	No	No	Yes	No	Yes
CBT	No	Yes	No	Yes	No	Yes
SSM	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv1	No	Yes	No	Yes	Yes, maybe	Yes, initially

Table 1: Multicast Routing Protocols Compared (Continued)

Multicast Routing Protocol	Dense Mode	Sparse Mode	Implicit Join	Explicit Join	(S,G) SBT	(* ,G) Shared Tree
IGMPv2	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv3	No	Yes	No	Yes	Yes, maybe	Yes, initially
BSR and Auto-RP	No	Yes	No	Yes	Yes, maybe	Yes, initially
MSDP	No	Yes	No	Yes	Yes, maybe	Yes, initially

It is important to realize that retransmissions due to a high bit-error rate on a link or overloaded routing device can make multicast as inefficient as repeated unicast. Therefore, there is a trade-off in many multicast applications regarding the session support provided by the Transmission Control Protocol (TCP) (but TCP always resends missing segments), or the simple drop-and-continue strategy of the User Datagram Protocol (UDP) datagram service (but reordering can become an issue). Modern multicast uses UDP almost exclusively.

T Series Router Multicast Performance

The Juniper Networks T Series Core Routers handle extreme multicast packet replication requirements with a minimum of router load. Each memory component replicates a multicast packet twice at most. Even in the worst-case scenario involving maximum fan-out, when 1 input port and 63 output ports need a copy of the packet, the T Series routing platform copies a multicast packet only six times. Most multicast distribution trees are much sparser, so in many cases only two or three replications are necessary. In no case does the T Series architecture have an impact on multicast performance, even with the largest multicast fan-out requirements.

Understanding Layer 3 Multicast Functionality on the SRX5K-MPC

Multicast is a “one source, many destinations” method of traffic distribution, meaning that only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

In the data plane of the SRX Series chassis, the SRX5000 line Module Port Concentrator (SRX5K-MPC) forwards Layer 3 IP multicast data packets, which include multicast protocol packets (for example, MLD, IGMP and PIM packets), and the data packets.

In incoming direction, the MPC receives multicast packets from an interface and forwards them to the central point or to a Services Processing Unit (SPU). The SPU performs multicast route lookup, flow-based security check, and packet replication.

In outgoing direction, the MPC receives copies of a multicast packet or Layer 3 multicast control protocol packets from SPU, and transmits them to either multicast capable routers or to hosts in a multicast group.

In the SRX Series chassis, the SPU perform multicast route lookup, if available, to forward an incoming multicast packet and replicates it for each multicast outgoing interface. After receiving replicated multicast packets and their corresponding outgoing interface information from the SPU, the MPC transmits these packets to next hops.



NOTE: On all SRX Series Firewalls, during RG1 failover with multicast traffic and high number of multicast sessions, the failover delay is from 90 through 120 seconds for traffic to resume on the secondary node. The delay of 90 through 120 seconds is only for the first failover. For subsequent failovers, the traffic resumes within 8 through 18 seconds.

RELATED DOCUMENTATION

| [Enabling PIM Sparse Mode](#) | 316

Multicast Configuration Overview

You configure a router network to support multicast applications with a related family of protocols. To use multicast, you must understand the basic components of a multicast network and their relationships, and then configure the device to act as a node in the network.

To configure the device as a node in a multicast network:

1. Determine whether the router is directly attached to any multicast sources.
Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers.
If receivers are present, IGMP is needed.
3. Determine whether to use the sparse, dense, or sparse-dense mode of multicast operation.

Each mode has different configuration considerations.

4. Determine the address of the rendezvous point (RP) if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, bootstrap router (BSR), or auto-RP method.

See:

- ["Understanding Static RP" on page 342](#)
 - ["Understanding the PIM Bootstrap Router" on page 365](#)
 - ["Understanding PIM Auto-RP" on page 371](#)
6. Determine whether to configure multicast to use its own reverse-path forwarding (RPF) routing table when configuring PIM in sparse, dense, or sparse-dense modes.
See ["Understanding Multicast Reverse Path Forwarding" on page 1149](#)
 7. (Optional) Configure the SAP and SDP protocols to listen for multicast session announcements.
See ["Configuring the Session Announcement Protocol" on page 573](#).
 8. Configure IGMP.
See ["Configuring IGMP" on page 25](#).
 9. (Optional) Configure the PIM static RP.
See ["Configuring Static RP" on page 342](#).
 10. (Optional) Filter PIM register messages from unauthorized groups and sources.
See ["Example: Rejecting Incoming PIM Register Messages on RP Routers" on page 389](#) and ["Example: Stopping Outgoing PIM Register Messages on a Designated Router" on page 382](#).
 11. (Optional) Configure a PIM RPF routing table.
See ["Example: Configuring a PIM RPF Routing Table" on page 1157](#).

RELATED DOCUMENTATION

[Multicast Overview | 2](#)

[Verifying a Multicast Configuration](#)

IPv6 Multicast Flow

IN THIS SECTION

- [IPv6 Multicast Flow Overview | 19](#)

IPv6 Multicast Flow Overview

The IPv6 multicast flow adds or enhances the following features:

- IPv6 transit multicast which includes the following packet functions:
 - Normal packet handling
 - Fragment handling
 - Packet reordering
- Protocol-Independent Multicast version 6 (PIMv6) flow handling
- Other multicast routing protocols, such as Multicast Listener Discovery (MLD)

The structure and processing of IPv6 multicast data session are the same as those of IPv4. Each data session has the following:

- One template session
- Several leaf sessions.

The reverse path forwarding (RPF) check behavior for IPv6 is the same as that for IPv4. Incoming multicast data is accepted only if the RPF check succeeds. In an IPv6 multicast flow, incoming Multicast Listener Discovery (MLD) protocol packets are accepted only if MLD or PIM is enabled in the security zone for the incoming interface. Sessions for multicast protocol packets have a default timeout value of 300 seconds. This value cannot be configured. The null register packet is sent to rendezvous point (RP).

In IPv6 multicast flow, a multicast router has the following three roles:

- Designated router

This router receives the multicast packets, encapsulates them with unicast IP headers, and sends them for multicast flow.
- Intermediate router

There are two sessions for the packets, the control session, for the outer unicast packets, and the data session. The security policies are applied to the data session and the control session, is used for forwarding.

- Rendezvous point

The RP receives the unicast PIM register packet, separates the unicast header, and then forwards the inner multicast packet. The packets received by RP are sent to the pd interface for decapsulation and are later handled like normal multicast packets.

On a Services Processing Unit (SPU), the multicast session is created as a template session for matching the incoming packet's tuple. Leaf sessions are connected to the template session. On the Customer Premise Equipment (CPE), only the template session is created. Each CPE session carries the fan-out lists that are used for load-balanced distribution of multicast SPU sessions.



NOTE: IPv6 multicast uses the IPv4 multicast behavior for session distribution.

The network service access point identifier (nsapi) of the leaf session is set up on the multicast text traffic going into the tunnels, to point to the outgoing tunnel. The zone ID of the tunnel is used for policy lookup for the leaf session in the second stage. Multicast packets are unidirectional. Thus for multicast text session sent into the tunnels, forwarding sessions are not created.

When the multicast route ages out, the corresponding chain of multicast sessions is deleted. When the multicast route changes, then the corresponding chain of multicast sessions is deleted. This forces the next packet hitting the multicast route to take the first path and re-create the chain of sessions; the multicast route counter is not affected.



NOTE: The IPv6 multicast packet reorder approach is same as that for IPv4.

For the encapsulating router, the incoming packet is multicast, and the outgoing packet is unicast. For the intermediate router, the incoming packet is unicast, and the outgoing packet is unicast.

RELATED DOCUMENTATION

| [Multicast Protocols User Guide](#)

Supported IP Multicast Protocol Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for IP multicast protocols, including the Distance Vector Multicast Routing Protocol (DVMRP), Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Multicast Source Discovery Protocol (MSDP), Pragmatic General Multicast (PGM), Protocol Independent Multicast (PIM), Session Announcement Protocol (SAP), and Session Description Protocol (SDP).

- RFC 1112, *Host Extensions for IP Multicasting* (defines IGMP Version 1)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2327, *SDP: Session Description Protocol*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3376, *Internet Group Management Protocol, Version 3*
- RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*
- RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*
- RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification*
- RFC 4604, *Using IGMPv3 and MLDv2 for Source-Specific Multicast*
- RFC 4607, *Source-Specific Multicast for IP*
- RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*
- RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*
- RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

The scoping mechanism is not supported.

- RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*
- RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

Starting in Release 17.3R1, Junos OS provides support for Protocol Independent Multicast (PIM) resolve type-length-value (TLV) for multicast in seamless MPLS. This support allows PIM in environments where the core routers do not maintain external routes.

- RFC 6513, *Multicast in MPLS/BGP IP VPNs*

- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-raggarwa-l3vpn-bgp-mvpn-extranet-08.txt, *Extranet in BGP Multicast VPN (MVPN)*
- Internet draft draft-rosen-l3vpn-spmi-joins-mldp-03.txt, *MVPN: S-PMSI Join Extensions for mLDP-Created Tunnels*

The following RFCs and Internet drafts do not define standards, but provide information about multicast protocols and related technologies. The IETF classifies them variously as “Best Current Practice,” “Experimental,” or “Informational.”

- RFC 1075, *Distance Vector Multicast Routing Protocol*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- RFC 2365, *Administratively Scoped IP Multicast*
- RFC 2547, *BGP/MPLS VPNs*
- RFC 2974, *Session Announcement Protocol*
- RFC 3208, *PGM Reliable Transport Protocol Specification*
- RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
- RFC 3569, *An Overview of Source-Specific Multicast (SSM)*
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*
- RFC 3973, *Protocol Independent Multicast – Dense Mode (PIM-DM): Protocol Specification (Revised)*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- Internet draft draft-ietf-idmr-dvmrp-v3-11.txt, *Distance Vector Multicast Routing Protocol*
- Internet draft draft-ietf-mboned-ssm232-08.txt, *Source-Specific Protocol Independent Multicast in 232/8*
- Internet draft draft-ietf-mmusic-sap-00.txt, *SAP: Session Announcement Protocol*
- Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*

Only section 7, “Data MDT: Optimizing flooding,” is supported.

RELATED DOCUMENTATION

| [Accessing Standards Documents on the Internet](#)

Ingress and Egress Multicast Replication using Recycle Replication Model

In ingress and egress replication with recycle port, there are two passes for each multicast packet. In the first pass, each packet is replicated to each egress Packet Forwarding Engine in the system using the Egress Packet Forwarding Engine's recycle port. In the second pass, packets are replicated to the ports of the egress Packet Forwarding Engine.

With this replication model, replication for the ports of a Packet Forwarding Engine occurs in the same Packet Forwarding Engine. Post replication, multicast packet flows are as good as unicast packet flows. There are no separate queues for multicast traffic. Unicast Traffic is scheduled as usual by getting credit from egress ports.

Recycle port is an internal logical port designed for loopback between egress and ingress pipeline in the system. Recycle port is created during boot time only if the recycle based replication is enabled from CLI. Use `set ingress-egress-recycle` statement at the system packet-forwarding-options mcast-repl-type hierarchy level for ingress and egress multicast replication using Recycle replication model. If you don't select the replication type, ingress replication (ingress-only) is considered as the default replication type.

The overall Multicast traffic flow towards egress PFE is restricted by recycle port bandwidth. Recycle port bandwidth is fixed with default bandwidth set to 30Gbps.

2

PART

Managing Group Membership

[Configuring IGMP and MLD | 25](#)

[Configuring IGMP Snooping | 96](#)

[Configuring MLD Snooping | 177](#)

[Configuring Multicast VLAN Registration | 245](#)

Configuring IGMP and MLD

IN THIS CHAPTER

- [Configuring IGMP | 25](#)
- [Verifying the IGMP Version | 57](#)
- [Configuring MLD | 59](#)
- [Understanding Distributed IGMP | 90](#)
- [Enabling Distributed IGMP | 92](#)

Configuring IGMP

IN THIS SECTION

- [Understanding Group Membership Protocols | 26](#)
- [Understanding IGMP | 27](#)
- [Configuring IGMP | 29](#)
- [Enabling IGMP | 31](#)
- [Modifying the IGMP Host-Query Message Interval | 32](#)
- [Modifying the IGMP Query Response Interval | 33](#)
- [Specifying Immediate-Leave Host Removal for IGMP | 34](#)
- [Filtering Unwanted IGMP Reports at the IGMP Interface Level | 35](#)
- [Accepting IGMP Messages from Remote Subnetworks | 36](#)
- [Modifying the IGMP Last-Member Query Interval | 37](#)
- [Modifying the IGMP Robustness Variable | 38](#)
- [Limiting the Maximum IGMP Message Rate | 40](#)
- [Changing the IGMP Version | 40](#)
- [Enabling IGMP Static Group Membership | 41](#)

- Recording IGMP Join and Leave Events | 50
- Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces | 52
- Tracing IGMP Protocol Traffic | 54
- Disabling IGMP | 56
- IGMP and Nonstop Active Routing | 57

Understanding Group Membership Protocols

There is a big difference between the multicast protocols used between host and routing device and between the multicast routing devices themselves. Hosts on a given subnetwork need to inform their routing device only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its routing devices only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routing devices of their participation in a multicast group. Between adjacent routing devices, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a routing device to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the routing device sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the routing device that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routing devices:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the routing device, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the routing device, especially on older or smaller routing devices.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routing devices can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM)*.

- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a routing device to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any routing device attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the routing device.



CAUTION: On MX Series platforms, IGMPv2 and IGMPv3 can or cannot be configured together on the same interface, depending on the Junos OS release at your installation. Configuring both together can cause unexpected behavior in multicast traffic forwarding.

SEE ALSO

| [Configuring MLD | 59](#)

Understanding IGMP

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routing devices in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routing devices. Multicast routing devices use IGMP to learn, for each of their attached physical networks, which groups have members.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

A routing device receives explicit join and prune messages from those neighboring routing devices that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The routing device then forwards data packets addressed to a multicast group G to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routing devices are automatically or statically designated as the RP, and all routing devices must explicitly join through the RP.

4. Each routing device along the path toward the RP builds a wildcard (any-source) state for the group and sends join and prune messages toward the RP.

The term *route entry* is used to refer to the state maintained in a routing device to represent the distribution tree.

A route entry can include such fields as:

- source address
- group address
- incoming interface from which packets are accepted
- list of outgoing interfaces to which packets are sent
- timers
- flag bits

The wildcard route entry's incoming interface points toward the RP.

The outgoing interfaces point to the neighboring downstream routing devices that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.

5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

Starting in Junos OS Release 15.2, PIMv1 is not supported.

IGMP is an integral part of IP and must be enabled on all routing devices and hosts that need to receive IP multicast traffic.

For each attached network, a multicast routing device can be either a querier or a nonquerier. The querier routing device periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a routing device can specify particular routing devices from which it accepts or rejects traffic. With IGMPv3, a multicast routing device can learn which sources are of interest to neighboring routing devices.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routing devices, IGMPv3 routing devices must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

SEE ALSO

[Supported IP Multicast Protocol Standards | 21](#)

[Enabling IGMP | 31](#)

[Disabling IGMP | 56](#)

[Configuring IGMP | 25](#)

Configuring IGMP

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See ["Configuring the Session Announcement Protocol" on page 573](#).

To configure the Internet Group Management Protocol (IGMP), include the `igmp` statement:

```
igmp {  
    accounting;  
    interface interface-name {  
        disable;  
    }  
}
```

```

(accounting | no-accounting);
group-policy [ policy-names ];
immediate-leave;
oif-map map-name;
promiscuous-mode;
ssm-map ssm-map-name;
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
version version;
}
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
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 protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).



NOTE: You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one “pseudo PIM interface” is created to represent all IGMP downstream (IGMP-only) interfaces on the

router. This reduces the amount of router resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

```
[edit protocols igmp]
user@host# set interface fe-0/0/0.0
```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0 {
    disable;
}
```

3. Enable IGMP on the interface by deleting the disable statement.

```
[edit protocols igmp]
delete interface ge-1/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0;
```

5. Verify the operation of IGMP on the interfaces by checking the output of the `show igmp interface` command.

SEE ALSO

[Understanding IGMP | 27](#)

[Disabling IGMP | 56](#)

show igmp interface

Modifying the IGMP Host-Query Message Interval

The objective of IGMP is to keep routers up to date with group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The IGMP querier router periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the `show igmp interface` command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the `show igmp statistics` command.

SEE ALSO

[Understanding IGMP | 27](#)

[Modifying the IGMP Query Response Interval | 33](#)

[Modifying the IGMP Robustness Variable | 38](#)

`show igmp interface`

`show igmp statistics`

Modifying the IGMP Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-response-interval 0.4
```

2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the `show igmp interface` command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the `show igmp statistics` command.

SEE ALSO

[Understanding IGMP | 27](#)

[Modifying the IGMP Host-Query Message Interval | 32](#)

[Modifying the IGMP Robustness Variable | 38](#)

show igmp interface

show igmp statistics

Specifying Immediate-Leave Host Removal for IGMP

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.



NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes

with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

```
[edit protocols IGMP]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the Immediate Leave field in the output of the `show igmp interface` command.

SEE ALSO

[Understanding IGMP | 27](#)

`show igmp interface`

Filtering Unwanted IGMP Reports at the IGMP Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The `group-policy` statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's `route-filter` statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's `route-filter` statement to match the group address and the policy's `source-address-filter` statement to match the source address.



CAUTION: On MX Series platforms, IGMPv2 and IGMPv3 can or cannot be configured together on the same interface, depending on the Junos OS release at your installation. Configuring both together can cause unexpected behavior in multicast traffic forwarding.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter 233.252.0.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject
```

2. Configure an IGMPv3 policy.

```
[edit policy-statement reject_policy_v3]
user@host# set from route-filter 233.252.0.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3
```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the `show igmp statistics` command.

SEE ALSO

[Understanding IGMP | 27](#)

Example: Configuring Policy Chains and Route Filters

show igmp statistics

Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the `promiscuous-mode` statement enables the routing device to accept IGMP messages from indirectly connected subnets.



NOTE: When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.



NOTE: When enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/1/1.0 promiscuous-mode
```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the `show igmp interface` command.
3. Verify the operation of the filter by checking the Rx non-local field in the output of the `show igmp statistics` command.

SEE ALSO

[Understanding IGMP | 27](#)

[Loopback Interface Configuration](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

`show igmp interface`

`show igmp statistics`

Modifying the IGMP Last-Member Query Interval

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member

query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols igmp]
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the IGMP Last Member Query Interval field in the output of the `show igmp interfaces` command.



NOTE: You can configure the last-member query count by configuring the robustness variable. The two are always equal.

SEE ALSO

[Modifying the IGMP Robustness Variable | 38](#)

`show pim interfaces`

Modifying the IGMP Robustness Variable

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query router receives an IGMP leave message on a shared network running IGMPv2, the query router must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- **Group member interval**—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- **Other querier present interval**—The robust count is used to calculate the amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- **Last-member query count**—Number of group-specific queries sent before the router assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

```
[edit protocols igmp]
user@host# set robust-count 5
```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the `show igmp interfaces` command.

SEE ALSO

[Modifying the IGMP Host-Query Message Interval | 32](#)

[Modifying the IGMP Query Response Interval | 33](#)

[Modifying the IGMP Last-Member Query Interval | 37](#)

`show pim interfaces`

Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the router.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a router with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the router can transmit in 1 second, include the `maximum-transmit-rate` statement and specify the maximum number of packets per second to be transmitted.

SEE ALSO

| *maximum-transmit-rate (Protocols IGMP)*

Changing the IGMP Version

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.



BEST PRACTICE: If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the `interface all` statement. That is, the new interface does not inherit the version number that you specified with the `interface all` statement. By default, that new interface is enabled with version 2. You must explicitly specify a `version number` when adding a new interface. For example, if you specified version 3 with `interface all`, you would need to configure the `version 3` statement for the new interface. Additionally, if you configure an interface for a multicast group at the `[edit interface interface-name static group multicast-group-address]` hierarchy level, you must specify a `version number` as well as the other group parameters. Otherwise, the interface is enabled with the default version 2.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0 version 3
```

2. Verify the configuration by checking the version field in the output of the `show igmp interfaces` command. The `show igmp statistics` command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.



CAUTION: On MX Series platforms, IGMPv2 and IGMPv3 can or cannot be configured together on the same interface, depending on the Junos OS release at your installation. Configuring both together can cause unexpected behavior in multicast traffic forwarding.

SEE ALSO

[Understanding IGMP | 27](#)

`show pim interfaces`

`show igmp statistics`

Enabling IGMP Static Group Membership

You can create IGMP static group membership to test multicast forwarding without a receiver host.

When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts. The router on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 233.252.0.1.

1. On the DR, configure the static groups to be created by including the static statement and group statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1
```

2. After you commit the configuration, use the show configuration protocol igmp command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  static {
    group 233.252.0.1 ;
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the show igmp group command to verify that static group 233.252.0.1 has been created.

```
user@host> show igmp group
Interface: fe-0/1/2
  Group: 233.252.0.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
```



NOTE: When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the `group-count` statement and specifying the number of groups to be created.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1 group-count 3
```

2. After you commit the configuration, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  static {
    group 233.252.0.1 {
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the `show igmp group` command to verify that static groups 233.252.0.1, 233.252.0.2, and 233.252.0.3 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
  Group: 233.252.0.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 233.252.0.2
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 233.252.0.3
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the `group-increment` statement and specifying the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1 group-count 3 group-increment
0.0.0.2
```

2. After you commit the configuration, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  version 3;
  static {
    group 233.252.0.1 {
      group-increment 0.0.0.2;
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the `show igmp group` command to verify that static groups 233.252.0.1, 233.252.0.3, and 233.252.0.5 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 233.252.0.1
```

```

Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
Group: 233.252.0.3
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
Group: 233.252.0.5
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 233.252.0.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the source statement and specifying the IPv4 address of the source host.

```

[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1 source 10.0.0.2

```

2. After you commit the configuration, use the show configuration protocol igmp command to verify the IGMP protocol configuration.

```

user@host> show configuration protocol igmp

```

```

interface fe-0/1/2.0 {
  version 3;
  static {
    group 233.252.0.1 {
      source 10.0.0.2;
    }
  }
}

```

```

    }
}

```

3. After you have committed the configuration and the source is sending traffic, use the `show igmp group` command to verify that static group 233.252.0.1 has been created and that source 10.0.0.2 has been accepted.

```

user@host> show igmp group
Interface: fe-0/1/2
  Group: 233.252.0.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 233.252.0.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the `source-count` statement and specifying the number of sources to be accepted.

```

[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1 source 10.0.0.2 source-count 3

```

2. After you commit the configuration, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```

user@host> show configuration protocol igmp

```

```

interface fe-0/1/2.0 {
  version 3;
  static {
    group 233.252.0.1 {
      source 10.0.0.2 {
        source-count 3;

```



```

    }
  }
}
}

```

3. After you have committed the configuration and the source is sending traffic, use the `show igmp group` command to verify that static group 233.252.0.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

```

user@host> show igmp group
Interface: fe-0/1/2
  Group: 233.252.0.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 233.252.0.1
    Source: 10.0.0.3
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 233.252.0.1
    Source: 10.0.0.4
    Last reported by: Local
    Timeout: 0 Type: Static

```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 233.252.0.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the `source-increment` statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```

[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1 source 10.0.0.2 source-count 3
source-increment 0.0.0.2

```

2. After you commit the configuration, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  version 3;
  static {
    group 233.252.0.1 {
      source 10.0.0.2 {
        source-count 3;
        source-increment 0.0.0.2;
      }
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the `show igmp group` command to verify that static group 233.252.0.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
  Group: 233.252.0.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 233.252.0.1
    Source: 10.0.0.4
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 233.252.0.1
    Source: 10.0.0.6
    Last reported by: Local
    Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 233.252.0.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the `exclude` statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 233.252.0.1 exclude source 10.0.0.2
```

2. After you commit the configuration, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {
  version 3;
  static {
    group 233.252.0.1 {
      exclude;
      source 10.0.0.2;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show igmp group detail` command to verify that static group 233.252.0.1 has been created and that the static group is operating in exclude mode.

```
user@host> show igmp group detail
Interface: fe-0/1/2
  Group: 233.252.0.1
    Group mode: Exclude
    Source: 10.0.0.2
```

Last reported by: Local
 Timeout: 0 Type: Static

SEE ALSO

[Enabling MLD Static Group Membership | 74](#)

group (Protocols IGMP)

group-count (Protocols IGMP)

group-increment (Protocols IGMP)

source-count (Protocols IGMP)

source-increment (Protocols IGMP)

static (Protocols IGMP)

Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

[Table 2 on page 50](#) describes the recordable IGMP events.

Table 2: IGMP Event Messages

ERRMSG Tag	Definition
RPD_IGMP_JOIN	Records IGMP join events.
RPD_IGMP_LEAVE	Records IGMP leave events.
RPD_IGMP_ACCOUNTING_ON	Records when IGMP accounting is enabled on an IGMP interface.
RPD_IGMP_ACCOUNTING_OFF	Records when IGMP accounting is disabled on an IGMP interface.
RPD_IGMP_MEMBERSHIP_TIMEOUT	Records IGMP membership timeout events.

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as **igmp-events**.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match “.*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* | .*RPD_IGMP_ACCOUNTING.*
| .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*”
```

3. Periodically archive the log file.

This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites “ftp://user@host1//var/tmp” password “anonymous”
user@host# set archive archive-sites “ftp://user@host2//var/tmp” password “test”
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start igmp-events
```

```
*** igmp-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command 'run monitor
start igmp-events '
monitor
```

SEE ALSO

[Understanding IGMP | 27](#)

Specifying Log File Size, Number, and Archiving Properties

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The `group-limit` statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (*,G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Starting in Junos OS Release 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The `group-threshold` statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the `group-limit` statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning

threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the `log-interval` statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.



NOTE: On ACX Series routers, the maximum number of multicast routes is 1024.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

```
[edit]
user@host# edit protocols igmp interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols igmp interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols igmp interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols igmp interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the `show protocols igmp` command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the `show igmp interface` command.

SEE ALSO

| [Enabling IGMP Static Group Membership](#) | 41

Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
client-notification	Trace notifications.
general	Trace general flow.
group	Trace group operations.
host-notification	Trace host notifications.
leave	Trace leave group messages (IGMPv2 only).
mtrace	Trace mtrace packets. Use the <code>mtrace</code> command to troubleshoot the software.
normal	Trace normal events.
packets	Trace all IGMP packets.
policy	Trace policy processing.
query	Trace IGMP membership query messages, including general and group-specific queries.
report	Trace membership report messages.
route	Trace routing information.

(Continued)

Flag	Description
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the IGMP trace file.

```
[edit protocols igmp traceoptions]
user@host# set file igmp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols igmp traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 233.252.0.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

SEE ALSO

[Understanding IGMP | 27](#)

[Tracing and Logging Junos OS Operations](#)

mtrace

Disabling IGMP

To disable IGMP on an interface, include the disable statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols igmp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*]



NOTE: ACX Series routers do not support [edit logical-systems *logical-system-name* protocols] hierarchy level.

SEE ALSO

[Understanding IGMP | 27](#)

[Configuring IGMP | 29](#)

[Enabling IGMP | 31](#)

IGMP and Nonstop Active Routing

Nonstop active routing (NSR) configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. These NSR configurations include passive support with IGMP in connection with PIM. The primary Routing Engine uses IGMP to determine its PIM multicast state, and this IGMP-derived information is replicated on the backup Routing Engine. IGMP on the new primary Routing Engine (after failover) relearns the state information quickly through IGMP operation. In the interim, the new primary Routing Engine retains the IGMP-derived PIM state as received by the replication process from the old primary Routing Engine. This state information times out unless refreshed by IGMP on the new primary Routing Engine. No additional IGMP configuration is required.

SEE ALSO

[Understanding Nonstop Active Routing for PIM | 515](#)

[Configuring MLD | 59](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.2	Starting in Junos OS Release 15.2, PIMv1 is not supported.
12.2	Starting in Junos OS Release 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface.

RELATED DOCUMENTATION

[Configuring MLD | 59](#)

Verifying the IGMP Version

IN THIS SECTION

● [Purpose | 58](#)

- Action | 58
- Meaning | 58

Purpose

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

Action

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

Sample Output

command-name

```
user@host> show igmp interface
Interface: ge-0/0/0.0
  Querier: 192.168.4.36
  State:      Up Timeout: 197 Version: 2 Groups: 0

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

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

Meaning

The output shows a list of the interfaces that are configured for IGMP. Verify the following information:

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

Configuring MLD

IN THIS SECTION

- [Understanding MLD | 59](#)
- [Configuring MLD | 63](#)
- [Enabling MLD | 64](#)
- [Modifying the MLD Version | 65](#)
- [Modifying the MLD Host-Query Message Interval | 66](#)
- [Modifying the MLD Query Response Interval | 67](#)
- [Modifying the MLD Last-Member Query Interval | 68](#)
- [Specifying Immediate-Leave Host Removal for MLD | 69](#)
- [Filtering Unwanted MLD Reports at the MLD Interface Level | 70](#)
- [Example: Modifying the MLD Robustness Variable | 71](#)
- [Limiting the Maximum MLD Message Rate | 73](#)
- [Enabling MLD Static Group Membership | 74](#)
- [Example: Recording MLD Join and Leave Events | 84](#)
- [Configuring the Number of MLD Multicast Group Joins on Logical Interfaces | 87](#)
- [Disabling MLD | 89](#)

Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 3 on page 60](#)). The querier routing device on the right is the receiver's DR.

Figure 3: Routing Devices Start Up on a Subnet




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In [Figure 4 on page 61](#), the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




NOTE: In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

Figure 4: Querier Routing Device Is Determined



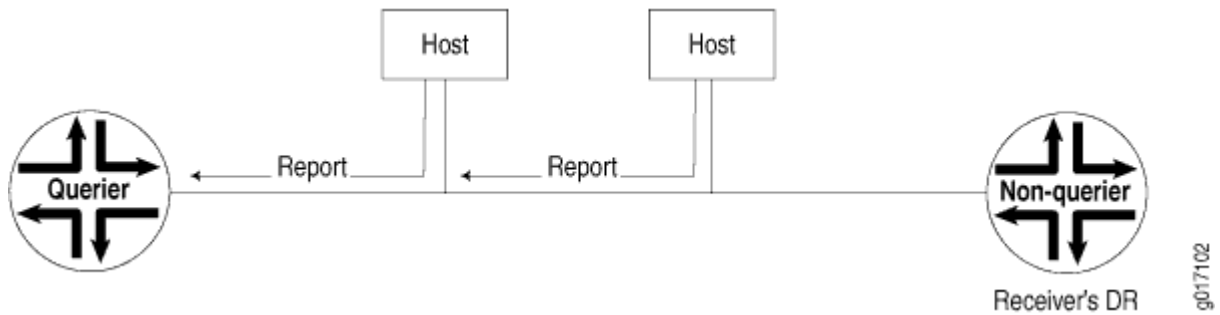
The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see [Figure 5 on page 61](#)). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 5: General Query Message Is Issued




If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 6 on page 62](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

Figure 6: Reports Are Received by the Querier Routing Device



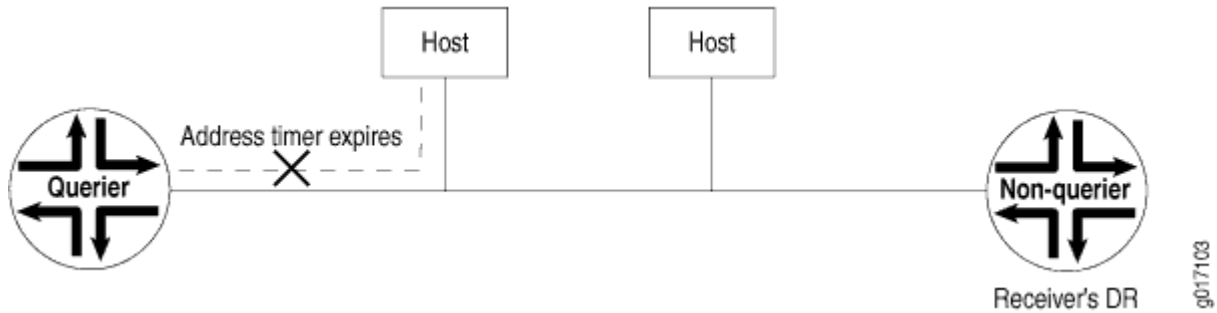
If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 7 on page 62](#)).

Figure 7: Host Has No Interested Receivers and Sends a Done Message to Routing Device



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 8 on page 63](#)).

Figure 8: Host Address Timer Expires and Address Is Removed from Multicast Address List



SEE ALSO

[Example: Recording MLD Join and Leave Events | 84](#)

[Example: Modifying the MLD Robustness Variable | 71](#)

Configuring MLD

To configure the Multicast Listener Discovery (MLD) Protocol, include the `mld` statement:

```
mld {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-policy [ policy-names ];
    immediate-leave;
    oif-map [ map-names ];
    passive;
    ssm-map ssm-map-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
  }
}
```

```

    version version;
  }
  maximum-transmit-rate packets-per-second;
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
}

```

You can include this statement at the following hierarchy levels:

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

By default, MLD is enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or the Distance Vector Multicast Routing Protocol (DVMRP).

Enabling MLD

The Multicast Listener Discovery (MLD) Protocol manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use MLD to learn which groups have members on each of their attached physical networks. MLD must be enabled for the router to receive IPv6 multicast packets. MLD is only needed for IPv6 networks, because multicast is handled differently in IPv4 networks. MLD is enabled on all IPv6 interfaces on which you configure PIM and on all IPv6 broadcast interfaces when you configure DVMRP.

MLD specifies different behaviors for multicast listeners and for routers. When a router is also a listener, the router responds to its own messages. If a router has more than one interface to the same link, it needs to perform the router behavior over only one of those interfaces. Listeners, on the other hand, must perform the listener behavior on all interfaces connected to potential receivers of multicast traffic.

If MLD is not running on an interface—either because PIM and DVMRP are not configured on the interface or because MLD is explicitly disabled on the interface—you can explicitly enable MLD.

To explicitly enable MLD:

1. If PIM and DVMRP are not running on the interface, explicitly enable MLD by including the interface name.

```

[edit protocols mld]
user@host# set interface fe-0/0/0.0

```

2. Check to see if MLD is disabled on any interfaces. In the following example, MLD is disabled on a Gigabit Ethernet interface.

```
[edit protocols mld]
user@host# show
```

```
interface fe-0/0/0.0;
interface ge-0/0/0.0 {
    disable;
}
```

3. Enable MLD on the interface by deleting the `disable` statement.

```
[edit protocols mld]
delete interface ge-0/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols mld]
user@host# show
```

```
interface fe-0/0/0.0;
interface ge-0/0/0.0;
```

5. Verify the operation of MLD by checking the output of the `show mld interface` command.

SEE ALSO

show mld interface

[CLI Explorer](#)

Modifying the MLD Version

By default, the router supports MLD version 1 (MLDv1). To enable the router to use MLD version 2 (MLDv2) for source-specific multicast (SSM) only, include the `version 2` statement.

If you configure the MLD version setting at the individual interface hierarchy level, it overrides configuring the IGMP version using the `interface all` statement.

If a source address is specified in a multicast group that is statically configured, the version must be set to MLDv2.

To change an MLD interface to version 2:

1. Configure the MLD interface.

```
[edit protocols mld]
user@host# set interface fe-0/0/0.0 version 2
```

2. Verify the configuration by checking the **version** field in the output of the `show mld interface` command. The `show mld statistics` command has version-specific output fields, such as the counters in the **MLD Message type** field.

SEE ALSO

[Source-Specific Multicast Groups Overview | 438](#)

[Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 458](#)

[Example: Configuring an SSM-Only Domain | 454](#)

[Example: Configuring PIM SSM on a Network | 451](#)

[Example: Configuring SSM Mapping | 455](#)

Modifying the MLD Host-Query Message Interval

The objective of MLD is to keep routers up to date with IPv6 group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The MLD querier router periodically sends general host-query messages on each attached network to solicit membership information. These messages solicit group membership information and are sent to the **link-scope all-nodes** address **FF02::1**. A general host-query message has a maximum response time that you can set by configuring the query response interval.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of MLD messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols mld]  
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the **MLD Query Interval** field in the output of the `show mld interface` command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the `show mld statistics` command.

SEE ALSO

show mld interface

[CLI Explorer](#)

show mld statistics

Modifying the MLD Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. You can change this interval to adjust the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

```
[edit protocols mld]
user@host# set query-response-interval 0.5
```

2. Verify the configuration by checking the **MLD Query Response Interval** field in the output of the `show mld interface` command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the `show mld statistics` command.

SEE ALSO

show mld interface

[CLI Explorer](#)

show mld statistics

Modifying the MLD Last-Member Query Interval

The last-member query interval (also called the last-listener query interval) is the maximum amount of time between group-specific query messages, including those sent in response to done messages sent on the **link-scope-all-routers** address FF02::2. You can lower this interval to reduce the amount of time it takes a router to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group (done) message from a host, the routing device sends multiple group-specific queries to the group. The querier sends a specific number of these queries, and it sends them at a specific interval. The number of queries sent is called the last-listener query count. The interval at which the queries are sent is called the last-listener query interval. Both settings are configurable, thus allowing you to adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-listener query count x (times) the last-listener query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-listener query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols mld]  
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the **MLD Last Member Query Interval** field in the output of the `show igmp interfaces` command.



NOTE: You can configure the last-member query count by configuring the robustness variable. The two are always equal.

SEE ALSO

show mld interface

[CLI Explorer](#)

Specifying Immediate-Leave Host Removal for MLD

The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.



NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host

on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave:

1. Configure immediate leave on the MLD interface.

```
[edit protocols mld]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the **Immediate Leave** field in the output of the `show mld interface` command.

Filtering Unwanted MLD Reports at the MLD Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The `group-policy` statement enables you to filter unwanted MLD reports at the interface level.

When the `group-policy` statement is enabled on a router, after the router receives an MLD report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only MLD group addresses (for MLDv1) by using the policy's `route-filter` statement to match the group address. You define the policy to match MLD (source, group) addresses (for MLDv2) by using the policy's `route-filter` statement to match the group address and the policy's `source-address-filter` statement to match the source address.

To filter unwanted MLD reports:

1. Configure an MLDv1 policy.

```
[edit policy-statement reject_policy_v1]
user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set then reject
```

2. Configure an MLDv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter fec0:1:1:4::/64 exact
```



```
user@host# set from source-address-filter fe80::2e0:81ff:fe05:1a8d/32 orlonger
user@host# set then reject
```

3. Apply the policies to the MLD interfaces where you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running MLDv1 and **ge-0/1/1.0** is running MLDv2.

```
[edit protocols mld]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v1
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v2
```

4. Verify the operation of the filter by checking the **Rejected Report** field in the output of the `show mld statistics` command.

SEE ALSO

[Routing Policies, Firewall Filters, and Traffic Policers User Guide](#)

`show mld statistics`

Example: Modifying the MLD Robustness Variable

IN THIS SECTION

- [Requirements | 71](#)
- [Overview | 72](#)
- [Configuration | 72](#)
- [Verification | 73](#)

This example shows how to configure and verify the MLD robustness variable in a multicast domain.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

- Enable IPv6 unicast routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Enable PIM. See "[PIM Overview](#)" on page 275.

Overview

The MLD robustness variable can be fine-tuned to allow for expected packet loss on a subnet. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

The value of the robustness variable is used in calculating the following MLD message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: $(\text{robustness variable} \times \text{query-interval}) + (1 \times \text{query-response-interval})$.
- Other querier present interval—Amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: $(\text{robustness variable} \times \text{query-interval}) + (0.5 \times \text{query-response-interval})$.
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The default number is the value of the robustness variable.

By default, the robustness variable is set to 2. The number can be from 2 through 10. You might want to increase this value if you expect a subnet to lose packets.

Configuration

IN THIS SECTION

- [Procedure | 73](#)

Procedure

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.

```
set protocols mld robust-count 5
```

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 [Junos OS CLI User Guide](#).

To change the value of the robustness variable:

1. Configure the robust count.

```
[edit protocols mld]  
user@host# set robust-count 5
```

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

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

Verification

To verify the configuration is working properly, check the **MLD Robustness Count** field in the output of the `show mld interfaces` command.

Limiting the Maximum MLD Message Rate

You can change the limit for the maximum number of MLD packets transmitted in 1 second by the router.

Increasing the maximum number of MLD packets transmitted per second might be useful on a router with a large number of interfaces participating in MLD.

To change the limit for the maximum number of MLD packets the router can transmit in 1 second, include the `maximum-transmit-rate` statement and specify the maximum number of packets per second to be transmitted.

Enabling MLD Static Group Membership

IN THIS SECTION

- [Create a MLD Static Group Member | 74](#)
- [Automatically create static groups | 75](#)
- [Automatically increment group addresses | 77](#)
- [Specify multicast source address \(in SSM mode\) | 78](#)
- [Automatically specify multicast sources | 79](#)
- [Automatically increment source addresses | 81](#)
- [Exclude multicast source addresses \(in SSM mode\) | 82](#)

Create a MLD Static Group Member

You can create MLD static group membership to test multicast forwarding without a receiver host. When you enable MLD static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts.

Class-of-service (CoS) adjustment is not supported with MLD static group membership.

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify the number of static groups to be automatically created.

In this example, you create static group `ff0e::1:ff05:1a8d`.

1. Configure the static groups to be created by including the `static` statement and `group` statement and specifying which IPv6 multicast address of the group to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d
```

2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d;
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the `show mld group` command to verify that static group `ff0e::1:ff05:1a8d` has been created.

```
user@host> show mld group
Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```



NOTE: You must specify a unique address for each group.

Automatically create static groups

When you create MLD static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. Configure the number of static groups to be created by including the `group-count` statement and specifying the number of groups to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d group-count 3
```

2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show mld group` command to verify that static groups `ff0e::1:ff05:1a8d`, `ff0e::1:ff05:1a8e`, and `ff0e::1:ff05:1a8f` have been created.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8e
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8f
    Source: fe80::2e0:81ff:fe05:1a8d
```

```
Last reported by: Local
Timeout: 0 Type: Static
```

Automatically increment group addresses

When you configure static groups on an interface on which you want to receive multicast traffic and you specify the number of static groups to be automatically created, you can also configure the group address to be automatically incremented by some number of addresses.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. Configure the group address increment by including the `group-increment` statement and specifying the number by which the address should be incremented for each group. The increment is specified in a format similar to an IPv6 address.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d group-count 3 group-
increment ::2
```

2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      group-increment ::2;
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show mld group` command to verify that static groups `ff0e::1:ff05:1a8d`, `ff0e::1:ff05:1a8f`, and `ff0e::1:ff05:1a91` have been created.

```
user@host> show mld group
```

```

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8f
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a91
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static

```

Specify multicast source address (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify the multicast source address to be accepted.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you create group ff0e::1:ff05:1a8d and accept IPv6 address fe80::2e0:81ff:fe05:1a8d as the only source.

1. Configure the source address by including the source statement and specifying the IPv6 address of the source host.

```

[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d source
fe80::2e0:81ff:fe05:1a8d

```


2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show mld group` command to verify that static group `ff0e::1:ff05:1a8d` has been created and that source `fe80::2e0:81ff:fe05:1a8d` has been accepted.

```
user@host> show mld group
```

```
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
```

Automatically specify multicast sources

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify a number of multicast sources to be automatically accepted.

In this example, you create static group `ff0e::1:ff05:1a8d` and accept `fe80::2e0:81ff:fe05:1a8d`, `fe80::2e0:81ff:fe05:1a8e`, and `fe80::2e0:81ff:fe05:1a8f` as the source addresses.

1. Configure the number of multicast source addresses to be accepted by including the `source-count` statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d source
fe80::2e0:81ff:fe05:1a8d source-count 3
```

2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show mld group` command to verify that static group `ff0e::1:ff05:1a8d` has been created and that sources `fe80::2e0:81ff:fe05:1a8d`, `fe80::2e0:81ff:fe05:1a8e`, and `fe80::2e0:81ff:fe05:1a8f` have been accepted.

```
user@host> show mld group
```

```
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8e
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
```

Automatically increment source addresses

When you configure static groups on an interface on which you want to receive multicast traffic, and specify a number of multicast sources to be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted.

In this example, you create static group `ff0e::1:ff05:1a8d` and accept `fe80::2e0:81ff:fe05:1a8d`, `fe80::2e0:81ff:fe05:1a8f`, and `fe80::2e0:81ff:fe05:1a91` as the sources.

1. Configure the number of multicast source addresses to be accepted by including the `source-increment` statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d source
fe80::2e0:81ff:fe05:1a8d source-count 3 source-increment ::2
```

2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
        source-increment ::2;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show mld group` command to verify that static group `ff0e::1:ff05:1a8d` has been created and that sources `fe80::2e0:81ff:fe05:1a8d`, `fe80::2e0:81ff:fe05:1a8f`, and `fe80::2e0:81ff:fe05:1a91` have been accepted.

```
user@host> show mld group
```

```
Interface: fe-0/1/2
```

```

Group: ff0e::1:ff05:1a8d
  Source: fe80::2e0:81ff:fe05:1a8d
  Last reported by: Local
  Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e2::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a91
    Last reported by: Local
    Timeout: 0 Type: Static

Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8f
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a91
Last reported by: Local
Timeout: 0 Type: Static

```

Exclude multicast source addresses (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the configured source address. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the configured source address.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you exclude address fe80::2e0:81ff:fe05:1a8d as a source for group ff0e::1:ff05:1a8d.

1. Configure a multicast static group to operate in exclude mode by including the `exclude` statement and specifying which IPv6 source address to be excluded.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d exclude source
fe80::2e0:81ff:fe05:1a8d
```

2. After you commit the configuration, use the `show configuration protocol mld` command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      exclude;
      source fe80::2e0:81ff:fe05:1a8d;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the `show mld group detail` command to verify that static group ff0e::1:ff05:1a8d has been created and that the static group is operating in exclude mode.

```
user@host> show mld group detail
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Group mode: Exclude
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
```

Similar configuration is available for IPv4 multicast traffic using the IGMP protocol.

RELATED DOCUMENTATION

| [Enabling IGMP Static Group Membership](#) | 41

Example: Recording MLD Join and Leave Events

IN THIS SECTION

- [Requirements](#) | 84
- [Overview](#) | 84
- [Configuration](#) | 85
- [Verification](#) | 87

This example shows how to determine whether MLD tuning is needed in a network by configuring the routing device to record MLD join and leave events.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Enable IPv6 unicast routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Enable PIM. See "[PIM Overview](#)" on page 275.

Overview

[Table 3 on page 84](#) describes the recordable MLD join and leave events.

Table 3: MLD Event Messages

ERRMSG Tag	Definition
RPD_MLD_JOIN	Records MLD join events.

Table 3: MLD Event Messages (*Continued*)

ERRMSG Tag	Definition
RPD_MLD_LEAVE	Records MLD leave events.
RPD_MLD_ACCOUNTING_ON	Records when MLD accounting is enabled on an MLD interface.
RPD_MLD_ACCOUNTING_OFF	Records when MLD accounting is disabled on an MLD interface.
RPD_MLD_MEMBERSHIP_TIMEOUT	Records MLD membership timeout events.

Configuration

IN THIS SECTION

- [Procedure | 85](#)

Procedure

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.

```
set protocols mld interface fe-0/1/0.2 accounting
set system syslog file mld-events any info
set system syslog file mld-events match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.*
| .*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
set system syslog file mld-events archive size 100000
set system syslog file mld-events archive files 3
set system syslog file mld-events archive transfer-interval 1440
set system syslog file mld-events archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
```

```
set system syslog file mld-events archive archive-sites "ftp://user@host2//var/tmp" password
"test"
```

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 [Junos OS CLI User Guide](#).

To configure recording of MLD join and leave events:

1. Enable accounting globally or on an MLD interface. This example shows the interface configuration.

```
[edit protocols mld]
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded, and filter the events to a system log file with a descriptive filename, such as **mld-events**.

```
[edit system syslog file mld-events]
user@host# set any info
[edit system syslog file mld-events]
user@host# set match “.*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* | .*RPD_MLD_ACCOUNTING.*
| .*RPD_MLD_MEMBERSHIP_TIMEOUT.*”
```

3. Periodically archive the log file.

This example rotates the file every 24 hours (1440 minutes) when it reaches 100 KB and keeps three files.

```
[edit system syslog file mld-events]
user@host# set archive size 100000
[edit system syslog file mld-events]
user@host# set archive files 3
[edit system syslog file mld-events]
user@host# set archive archive-sites “ftp://user@host1//var/tmp” password “anonymous”
[edit system syslog file mld-events]
user@host# set archive archive-sites “ftp://user@host2//var/tmp” password “test”
[edit system syslog file mld-events]
user@host# set archive transfer-interval 1440
```



```
[edit system syslog file mld-events]
user@host# set archive start-time 2011-01-07:12:30
```

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

```
[edit system syslog file mld-events]]
user@host# commit
```

Verification

You can view the system log file by running the **file show** command.

```
user@host> file show mld-events
```

You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start mld-events
```

```
*** mld-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command 'run monitor start
mld-events '
monitor
```

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

The `group-limit` statement enables you to limit the number of MLD multicast group joins for logical interfaces. When this statement is enabled on a router running MLD version 2, the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for MLD multicast groups, keep the following in mind:

- Each any-source group (*,G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in MLDv2 exclude mode are counted toward the limit.

- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on MLD logical interfaces by using dynamic profiles. For detailed information about creating dynamic profiles, see the [Junos OS Subscriber Management and Services Library](#).

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for MLD multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of MLD multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The `group-threshold` statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the `group-limit` statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs a warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the `log-interval` statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for MLD multicast group joins.

To limit multicast group joins on an MLD logical interface:

1. Access the logical interface at the MLD protocol hierarchy level.

[edit]

```
user@host# edit protocols mld interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols mld interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols mld interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols mld interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the `show protocols mld` command. To verify the operation of MLD on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the `show mld interface` command.

Disabling MLD

To disable MLD on an interface, include the `disable` statement:

```
interface interface-name {
  disable;
}
```

You can include this statement at the following hierarchy levels:

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

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
12.2	Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for MLD multicast group joins received on the logical interface.

RELATED DOCUMENTATION

| [Configuring IGMP | 25](#)

Understanding Distributed IGMP

IN THIS SECTION

- [Distributed IGMP Overview | 90](#)
- [Guidelines for Configuring Distributed IGMP | 91](#)

By default, Internet Group Management Protocol (IGMP) processing takes place on the Routing Engine for MX Series routers. This centralized architecture may lead to reduced performance in scaled environments or when the Routing Engine undergoes CLI changes or route updates. You can improve system performance for IGMP processing by enabling *distributed IGMP*, which utilizes the Packet Forwarding Engine to maintain a higher system-wide processing rate for join and leave events.

Distributed IGMP Overview

Distributed IGMP works by moving IGMP processing from the Routing Engine to the Packet Forwarding Engine. When distributed IGMP is not enabled, IGMP processing is centralized on the routing protocol process (rpd) running on the Routing Engine. When you enable distributed IGMP, join and leave events are processed across Modular Port Concentrators (MPCs) on the Packet Forwarding Engine. Because join and leave processing is distributed across multiple MPCs instead of being processed through a centralized rpd on the Routing Engine, performance improves and join and leave latency decreases.

When you enable distributed IGMP, each Packet Forwarding Engine processes reports and generates queries, maintains local group membership to the interface mapping table and updates the forwarding state based on this table, runs distributed IGMP independently, and implements the `group-policy` and `ssm-map-policy` IGMP interface options.



NOTE: Information from `group-policy` and `ssm-map-policy` IGMP interface options passes from the Routing Engine to the Packet Forwarding Engine.

When you enable distributed IGMP, the rpd on the Routing Engine synchronizes all IGMP configurations (including global and interface-level configurations) from the rpd to each Packet Forwarding Engine, runs

passive IGMP on distributed interfaces, and notifies Protocol Independent Multicast (PIM) of all group memberships per distributed IGMP interface.

Guidelines for Configuring Distributed IGMP

Consider the following guidelines when you configure distributed IGMP on an MX Series router with MPCs:

- Distributed IGMP increases network performance by reducing the maximum join and leave latency and by increasing join and leave events.



NOTE: Join and leave latency may increase if multicast traffic is not preprovisioned and destined for an MX Series router when a join or leave event is received from a client interface.

- Distributed IGMP is supported for Ethernet interfaces. It does not improve performance on PIM interfaces.
- Starting in Junos OS release 18.2, distributed IGMP is supported on Ethernet interfaces with enhanced subscriber management. IGMP processing for subscriber flows is moved from the Routing Engine to the Packet Forwarding Engine of supported line cards.

Multicast groups cannot be comprised of mixed receivers. They can be either centralized IGMP or distributed IGMP.

- You can reduce initial join delays by enabling Protocol Independent Multicast (PIM) static joins or IGMP static joins. You can reduce initial delays even more by preprovisioning multicast traffic. When you preprovision multicast traffic, MPCs with distributed IGMP interfaces receive multicast traffic.
- For distributed IGMP to function properly, you must enable enhanced IP network services on a single-chassis MX Series router. Virtual Chassis is not supported.
- When you enable distributed IGMP, the following interface options are not supported on the Packet Forwarding Engine: `oif-map`, `group-limit`, `ssm-map`, and `static`. The `traceoptions` and `accounting` statements can only be enabled for IGMP operations still performed on the Routing Engine; they are not supported on the Packet Forwarding Engine. The `clear igmp membership` command is not supported when distributed IGMP is enabled.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.2	Starting in Junos OS Release 18.2, distributed IGMP is supported on Ethernet interfaces with enhanced subscriber management. IGMP processing for subscriber flows is moved from the Routing Engine to the Packet Forwarding Engine of supported line cards. Multicast groups cannot be comprised of mixed receivers. They can be either centralized IGMP or distributed IGMP.

RELATED DOCUMENTATION

[Understanding IGMP | 27](#)

[Junos OS Multicast Protocols User Guide](#)

Enabling Distributed IGMP

IN THIS SECTION

- [Enabling Distributed IGMP on Static Interfaces | 92](#)
- [Enabling Distributed IGMP on Dynamic Interfaces | 93](#)
- [Configuring Multicast Traffic for Distributed IGMP | 93](#)

Configuring distributed IGMP improves performance by reducing join and leave latency. This works by moving IGMP processing from the Routing Engine to the Packet Forwarding Engine. In contrast to centralized IGMP processing on the Routing Engine, the Packet Forwarding Engine disperses traffic across multiple Modular Port Concentrators (MPCs).

You can enable distributed IGMP on static interfaces or dynamic interfaces. As a prerequisite, you must enable enhanced IP network services on a single-chassis MX Series router.

Enabling Distributed IGMP on Static Interfaces

You can enable distributed IGMP on a static interface by configuring enhanced IP network services and including the distributed statement at the [edit protocols igmp interface *interface-name*] hierarchy level. Enhanced IP network services must be enabled (at the [chassis network-services enhanced-ip] hierarchy).

To enable distributed IGMP on a static interface:

1. Configure the IGMP static interface.

```
[edit protocols igmp ]
user@host# set interface interface-name
```

2. Enable distributed IGMP on a static interface.

```
[edit protocols igmp interface interface-name]
user@host# set distributed
```

3. Commit the configuration.

Enabling Distributed IGMP on Dynamic Interfaces

You can enable distributed IGMP on a dynamic interface by configuring enhanced IP network services and including the distributed statement at the [edit dynamic profiles *profile-name* protocols] hierarchy level. Enhanced IP network services must be enabled (at the [chassis network-services enhanced-ip] hierarchy).

1. Configure the IGMP interface.

```
[edit dynamic profiles profile-name protocols]
user@host# set interface $junos-interface-name
```

2. Enable distributed IGMP on a dynamic interface.

```
[edit dynamic profiles profile-name protocols interface $junos-interface-name]
user@host# set distributed
```

3. Commit the configuration.

Configuring Multicast Traffic for Distributed IGMP

Configuring static source and group (S,G) addresses for distributed IGMP reduces join delays and sends multicast traffic to the last-hop router. You can configure static multicast groups (S,G) for distributed IGMP at the [edit protocols pim] hierarchy level. You can issue the distributed keyword at one of the following three hierarchy levels:

- [edit protocols pim static]

Issuing the `distributed` keyword at this hierarchy level enables static joins for specific multicast (S,G) groups and preprovisions all of them so that all distributed IGMP Packet Forwarding Engines receive traffic.

- [edit protocols pim static group *multicast-group-address*]

Issuing the `distributed` keyword at this hierarchy level enables static joins for multicast (S,G) groups so that all distributed IGMP Packet Forwarding Engines receive traffic and preprovisions a specific multicast group address (G).

- [edit protocols pim static group *multicast-group-address* source *source-address*]

Issuing the `distributed` keyword at this hierarchy level enables static joins for multicast (S,G) groups so that all Packet Forwarding Engines receive traffic, but preprovisions a specific multicast (S,G) group.

To configure static multicast (S,G) addresses for distributed IGMP:

1. Configure static PIM.

```
[edit protocols pim]
user@host# set static
```

2. (Optional) Enable static joins for specific (S,G) addresses and preprovision all of them so that all distributed IGMP Packet Forwarding Engines receive traffic. In the example, multicast traffic for all of the groups (225.0.0.1, 10.10.10.1), (225.0.0.1, 10.10.10.2), and (225.0.0.2, *) is preprovisioned.

```
[edit protocols pim]
user@host# set protocols pim static distributed
user@host# set protocols pim static group 225.0.0.1 source 10.10.10.1
user@host# set protocols pim static group 225.0.0.1 source 10.10.10.2
user@host# set protocols pim static group 225.0.0.2
```

3. (Optional) Enable static joins for specific multicast (S,G) groups so that all distributed IGMP Packet Forwarding Engines receive traffic and preprovision a specific multicast group address (G). In the example, multicast traffic for groups (225.0.0.1, 10.10.10.1) and (225.0.0.1, 10.10.10.2) is preprovisioned, but group (225.0.0.2, *) is not preprovisioned.

```
[edit protocols pim]
user@host# set protocols pim static
user@host# set protocols pim static group 225.0.0.1 distributed
user@host# set protocols pim static group 225.0.0.1 source 10.10.10.1
```



```
user@host# set protocols pim static group 225.0.0.1 source10.10.10.2
user@host# set protocols pim static group 225.0.0.2
```

4. (Optional) Enable a static join for specific multicast (S,G) groups so that all Packet Forwarding Engines receive traffic, but preprovision only one specific multicast address group. In the example, multicast traffic for group (225.0.0.1, 10.10.10.1) is preprovisioned, but all other groups are not preprovisioned.

```
[edit protocols pim]
user@host# set protocols pim static
user@host# set protocols pim static group 225.0.0.1
user@host# set protocols pim static group 225.0.0.1 source 10.10.10.1 distributed
user@host# set protocols pim static group225.0.0.1 source10.10.10.2
user@host# set protocols pim static group 225.0.0.2
```

5. Commit the configuration.

SEE ALSO

Configuring Dynamic DHCP Client Access to a Multicast Network

[Junos OS Multicast Protocols User Guide](#)

[Junos OS Multicast Protocols User Guide](#)

Configuring IGMP Snooping

IN THIS CHAPTER

- [IGMP Snooping Overview | 96](#)
- [Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment | 104](#)
- [Configuring IGMP Snooping on Switches | 122](#)
- [Example: Configuring IGMP Snooping on Switches | 126](#)
- [Example: Configuring IGMP Snooping on EX Series Switches | 130](#)
- [Verifying IGMP Snooping on EX Series Switches | 134](#)
- [Changing the IGMP Snooping Group Timeout Value on Switches | 138](#)
- [Monitoring IGMP Snooping | 139](#)
- [Example: Configuring IGMP Snooping | 141](#)
- [Example: Configuring IGMP Snooping on SRX Series Devices | 167](#)
- [Configuring Point-to-Multipoint LSP with IGMP Snooping | 174](#)

IGMP Snooping Overview

IN THIS SECTION

- [Benefits of IGMP Snooping | 97](#)
- [How IGMP Snooping Works | 97](#)
- [How IGMP Snooping Works with Routed VLAN Interfaces | 98](#)
- [IGMP Message Types | 98](#)
- [How Hosts Join and Leave Multicast Groups | 99](#)
- [Support for IGMPv3 Multicast Sources | 99](#)
- [IGMP Snooping and Forwarding Interfaces | 100](#)

- [General Forwarding Rules | 101](#)
- [Using the Device as an IGMP Querier | 101](#)
- [IGMP Snooping on Private VLANs \(PVLANS\) | 102](#)

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a device. With IGMP snooping enabled, the device monitors IGMP traffic on the network and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. The device conserves bandwidth by sending multicast traffic only to interfaces connected to devices that want to receive the traffic, instead of flooding the traffic to all the downstream interfaces in a VLAN.

Benefits of IGMP Snooping

- **Optimized bandwidth utilization**—IGMP snooping's main benefit is to reduce flooding of packets. The device selectively forwards IPv4 multicast data to a list of ports that want to receive the data instead of flooding it to all ports in a VLAN.
- **Improved security**—Prevents denial of service attacks from unknown sources.

How IGMP Snooping Works

Devices usually learn unicast MAC addresses by checking the source address field of the frames they receive and then send any traffic for that unicast address only to the appropriate interfaces. However, a multicast MAC address can never be the source address for a packet. As a result, when a device receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, sending a significant amount of traffic for which there might not necessarily be interested receivers.

IGMP snooping prevents this flooding. When you enable IGMP snooping, the device monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the device receives multicast packets, it uses the multicast forwarding table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups.

On EX Series and QFX Series switches that do not support the Enhanced Layer 2 Software (ELS) configuration style, IGMP snooping is enabled by default on all VLANs (or only on the default VLAN on some devices) and you can disable it selectively on one or more VLANs. On all other devices, you must explicitly configure IGMP snooping on a VLAN or in a bridge domain to enable it.



NOTE: You can't configure IGMP snooping on a secondary (private) VLAN (PVLAN). However, starting in Junos OS Release 18.3R1 on EX4300 switches and EX4300 Virtual Chassis, and Junos OS Release 19.2R1 on EX4300 multigigabit switches, when you enable IGMP snooping on a primary VLAN, you also implicitly enable it on any secondary VLANs defined for that primary VLAN. See ["IGMP Snooping on Private VLANs \(PVLANS\)" on page 102](#) for details.

How IGMP Snooping Works with Routed VLAN Interfaces

The device can use a routed VLAN interface (RVI) to forward traffic between VLANs in its configuration. IGMP snooping works with Layer 2 interfaces and RVIs to forward multicast traffic in a switched network.

When the device receives a multicast packet, its Packet Forwarding Engines perform a multicast lookup on the packet to determine how to forward the packet to its local interfaces. From the results of the lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces that have ports local to the Packet Forwarding Engine. If the list includes an RVI, the device provides a bridge multicast group ID for the RVI to the Packet Forwarding Engine.

For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID, which identifies the Layer 2 interfaces in the VLAN that are interested in receiving the multicast stream. The Packet Forwarding Engine then forwards multicast traffic to bridge multicast IDs that have multicast receivers for a given multicast group.

IGMP Message Types

Multicast routers use IGMP to learn which groups have interested listeners for each of their attached physical networks. In any given subnet, one multicast router acts as an IGMP querier. The IGMP querier sends out the following types of queries to hosts:

- General query—Asks whether any host is listening to any group.
- Group-specific query—(IGMPv2 and IGMPv3 only) Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.
- Group-and-source-specific query—(IGMPv3 only) Asks whether any host is listening to group multicast traffic from a specific multicast source. This query is sent in response to a host indicating that it is no longer interested in receiving group multicast traffic from the multicast source and allows the router to quickly determine any remaining hosts are interested in receiving group multicast traffic from that source.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report—Indicates that the host wants to join a particular multicast group.
- Leave report—(IGMPv2 and IGMPv3 only) Indicates that the host wants to leave a particular multicast group.

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group the host wants to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, it must continue to respond to the periodic general IGMP queries.

Hosts can leave a multicast group in either of two ways:

- By not responding to periodic queries within a particular interval of time, which is considered a “silent leave.” This is the only leave method for IGMPv1 hosts.
- By sending a leave report. This method can be used by IGMPv2 and IGMPv3 hosts.

Support for IGMPv3 Multicast Sources

In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source *except* the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources.

Devices that support IGMPv3 process INCLUDE and EXCLUDE membership reports, and most devices forward source-specific multicast (SSM) traffic only from requested sources to subscribed receivers accordingly. However, you might see the device doesn't strictly forward multicast traffic on a per-source basis in some configurations such as:

- EX Series and QFX Series switches that do not use the Enhanced Layer 2 Software (ELS) configuration style
- EX2300 and EX3400 switches running Junos OS Releases prior to 18.1R2

- EX4300 switches running Junos OS Releases prior to 18.2R1, 18.1R2, 17.4R2, 17.3R3, 17.2R3, and 14.1X53-D47
- SRX Series Services Gateways

In these cases, the device might consolidate all INCLUDE and EXCLUDE mode reports they receive on a VLAN for a specified group into a single route that includes all multicast sources for that group, with the next hop representing all interfaces that have interested receivers for the group. As a result, interested receivers on the VLAN can receive traffic from a source that they did not include in their INCLUDE report or from a source they excluded in their EXCLUDE report. For example, if Host 1 wants traffic for G from Source A and Host 2 wants traffic for group G from Source B, they both receive traffic for group G regardless of whether A or B sends the traffic.

IGMP Snooping and Forwarding Interfaces

To determine how to forward multicast traffic, the device with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The device learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the device adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the device adds the interface to its multicast forwarding table as a group-member interface.

Learned interface table entries age out after a time period. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the device removes the entry for that interface from its multicast forwarding table.



NOTE: For the device to learn multicast-router interfaces and group-member interfaces, the network must include an IGMP querier. This is often in a multicast router, but if there is no multicast router on the local network, you can configure the device itself to be an IGMP querier.

You can statically configure an interface to be a multicast-router interface or a group-member interface. The device adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. A device can have a mix of statically configured and dynamically learned interfaces.

General Forwarding Rules

An interface in a VLAN with IGMP snooping enabled receives multicast traffic and forwards it according to the following rules.

IGMP traffic:

- Forward IGMP general queries received on a multicast-router interface to all other interfaces in the VLAN.
- Forward IGMP group-specific queries received on a multicast-router interface to only those interfaces in the VLAN that are members of the group.
- Forward IGMP reports received on a host interface to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not IGMP traffic:

- Flood multicast packets with a destination address of 233.252.0.0/24 to all other interfaces on the VLAN.
- Forward unregistered multicast packets (packets for a group that has no current members) to all multicast-router interfaces in the VLAN.
- Forward registered multicast packets to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

Using the Device as an IGMP Querier

With IGMP snooping on a pure Layer 2 local network (that is, Layer 3 is not enabled on the network), if the network doesn't include a multicast router, multicast traffic might not be properly forwarded through the network. You might see this problem if the local network is configured such that multicast traffic must be forwarded between devices in order to reach a multicast receiver. In this case, an upstream device does not forward multicast traffic to a downstream device (and therefore to the multicast receivers attached to the downstream device) because the downstream device does not forward IGMP reports to the upstream device. You can solve this problem by configuring one of the devices to be an IGMP querier. The IGMP querier device sends periodic general query packets to all the devices in the network, which ensures that the snooping membership tables are updated and prevents multicast traffic loss.

If you configure multiple devices to be IGMP queriers, the device with the lowest (smallest) IGMP querier source address takes precedence and acts as the querier. The devices with higher IGMP querier source addresses stop sending IGMP queries unless they do not receive IGMP queries for 255 seconds. If the device with a higher IGMP querier source address does not receive any IGMP queries during that period, it starts sending queries again.



NOTE: QFabric systems in Junos OS Release 14.1X53-D15 support the `igmp-querier` statement, but do not support this statement in Junos OS 15.1.

To configure a device to act as an IGMP querier, enter the following:

```
[edit protocols]
user@host# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

To configure a QFabric Node device switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@host# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

IGMP Snooping on Private VLANs (PVLANS)

A PVLAN consists of secondary isolated and community VLANs configured within a primary VLAN. Without IGMP snooping support on the secondary VLANs, multicast streams received on the primary VLAN are flooded to the secondary VLANs.

Starting in Junos OS Release 18.3R1, EX4300 switches and EX4300 Virtual Chassis support IGMP snooping with PVLANS. Starting in Junos OS Release 19.2R1, EX4300 multigigabit model switches support IGMP snooping with PVLANS. When you enable IGMP snooping on a primary VLAN, you also implicitly enabled it on all secondary VLANs. The device learns and stores multicast group information on the primary VLAN, and also learns the multicast group information on the secondary VLANs in the context of the primary VLAN. As a result, the device further constrains multicast streams only to interested receivers on secondary VLANs, rather than flooding the traffic in all secondary VLANs.

The CLI prevents you from explicitly configuring IGMP snooping on secondary isolated or community VLANs. You only need to configure IGMP snooping on the primary VLAN under which the secondary VLANs are defined. For example, for a primary VLAN `vlan-pri` with a secondary isolated VLAN `vlan-iso` and a secondary community VLAN `vlan-comm`:

```
set vlans vlan-pri vlan-id 100
set vlans vlan-pri isolated-vlan vlan-iso
set vlans vlan-pri community-vlans vlan-comm
set vlans vlan-iso vlan-id 300
set vlans vlan-iso private-vlan isolated
set vlans vlan-comm vlan-id 200
```



```
set vlans vlan-comm private-vlan community
set protocols igmp-snooping vlan vlan-pri
```

IGMP reports and leave messages received on secondary VLAN ports are learned in the context of the primary VLAN. Promiscuous trunk ports or inter-switch links acting as multicast router interfaces for the PVLAN receive incoming multicast data streams from multicast sources and forward them only to the secondary VLAN ports with learned multicast group entries.

This feature does not support secondary VLAN ports as multicast router interfaces. The CLI does not strictly prevent you from statically configuring an interface on a community VLAN as a multicast router port, but IGMP snooping does not work properly on PVLANS with this configuration. When IGMP snooping is configured on a PVLAN, the switch also automatically disables dynamic multicast router port learning on any isolated or community VLAN interfaces. IGMP snooping with PVLANS also does not support configurations with an IGMP querier on isolated or community VLAN interfaces.

See [Understanding Private VLANs](#) and [Creating a Private VLAN Spanning Multiple EX Series Switches with ELS Support \(CLI Procedure\)](#) for details on configuring PVLANS.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.2R1	Starting in Junos OS Release 19.2R1, EX4300 multigigabit model switches support IGMP snooping with PVLANS.
18.3R1	Starting in Junos OS Release 18.3R1, EX4300 switches and EX4300 Virtual Chassis support IGMP snooping with PVLANS.
14.1X53-D15	QFabric systems in Junos OS Release 14.1X53-D15 support the <code>igmp-querier</code> statement, but do not support this statement in Junos OS 15.1.

RELATED DOCUMENTATION

[Example: Configuring IGMP Snooping on SRX Series Devices | 167](#)

[Example: Configuring IGMP Snooping on Switches | 126](#)

[Configuring IGMP Snooping on Switches | 122](#)

[Monitoring IGMP Snooping | 139](#)

[Configuring IGMP | 29](#)

Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment

IN THIS SECTION

- Benefits of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment | 106
- Supported IGMP or MLD Versions and Group Membership Report Modes | 106
- Summary of Multicast Traffic Forwarding and Routing Use Cases | 107
- Use Case 1: Intra-VLAN Multicast Traffic Forwarding | 108
- Use Case 2: Inter-VLAN Multicast Routing and Forwarding—IRB Interfaces with PIM | 111
- Use Case 3: Inter-VLAN Multicast Routing and Forwarding—PIM Gateway with Layer 2 Connectivity | 115
- Use Case 4: Inter-VLAN Multicast Routing and Forwarding—PIM Gateway with Layer 3 Connectivity | 117
- Use Case 5: Inter-VLAN Multicast Routing and Forwarding—External Multicast Router | 119
- EVPN Multicast Flags Extended Community | 119

Internet Group Management Protocol (IGMP) snooping and Multicast Listener Discovery (MLD) snooping constrain multicast traffic in a broadcast domain to interested receivers and multicast devices. In an environment with a significant volume of multicast traffic, using IGMP or MLD snooping preserves bandwidth because multicast traffic is forwarded only on those interfaces that have multicast listeners. IGMP snooping optimizes IPv4 multicast traffic flow. MLD snooping optimizes IPv6 multicast traffic flow.

We support IGMP snooping and MLD snooping in Ethernet VPN-Virtual extensible VLAN (EVPN-VXLAN) networks as a way to help optimize multicast traffic flow.

Starting in Junos OS Release 21.2R1, on some platforms we also support optimized intersubnet multicast (OISM) routing and forwarding in EVPN-VXLAN edge-routed bridging (ERB) overlay networks. See [Optimized Inter-Subnet Multicast in EVPN Networks](#) for full details on OISM configuration and operation. OISM enables efficient multicast routing and forwarding for both internal and external multicast sources and receivers in ERB overlay fabrics. You must configure IGMP snooping (or, if supported, MLD snooping) on the fabric leaf devices as part of OISM configuration. On some platforms, we support EVPN-VXLAN multicast traffic only with OISM.

This topic provides an introduction to available multicast optimization methods without OISM in EVPN-VXLAN CRB overlays or ERB overlays.

EVPN-VXLAN devices might support IGMPv2 and IGMPv3 with IGMP snooping for IPv4 multicast traffic. The devices might also support MLDv1 and MLDv2 with MLD snooping for IPv6 multicast traffic.

- With IGMPv2 and MLDv1, the device processes only any-source multicast (ASM) reports.
- If the device supports IGMPv3 and MLDv2, ASM mode is the default behavior, but you can configure the device to process source-specific multicast (SSM) reports with IGMPv3 and MLDv2 instead.

However, the device can't process both SSM reports and ASM reports at the same time. When you configure a device to operate in SSM mode with IGMPv3 and MLDv2, the device drops any ASM reports. When you don't configure the device to operate in SSM mode, a device processes any ASM reports but drops IGMPv3 and MLDv2 SSM reports.

Unless mentioned explicitly, the information in this topic applies to IGMPv2, IGMPv3, MLDv1, and MLDv2 on the devices that support these protocols in the following IP fabric architectures:

- EVPN-VXLAN ERB overlay
- EVPN-VXLAN CRB overlay



NOTE: On a switching devices, you can configure a *VLAN*. On routing devices, you can configure the same entity called a *bridge domain*. To keep things simple, this topic uses the term *VLAN* when we refer to the same entity configured on either switching or routing devices.

Here are a few important general notes about IGMP snooping and MLD snooping behavior with EVPN-VXLAN:



NOTE:

- On switches in the QFX5000 line, in releases up until Junos OS Releases 18.4R2 and 19.1R2, with IGMP snooping enabled the switches only constrain flooding for multicast traffic coming in on the VXLAN tunnel network ports. The switches would still flood multicast traffic coming in from an access interface to all other access and network interfaces.
- Starting with Junos release 22.3R1, on some platforms we support IPv4 and IPv6 multicast traffic in an IPv6 EVPN-VXLAN overlay with IPv6 underlay peering.

Benefits of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment

- In an environment with a significant volume of multicast traffic, using IGMP snooping or MLD snooping constrains multicast traffic in a VLAN to interested receivers and multicast devices, which conserves network bandwidth.
- Synchronizing the IGMP or MLD state among all EVPN devices for multihomed receivers ensures that all subscribed listeners receive multicast traffic, even in cases such as the following:
 - IGMP or MLD membership reports for a multicast group might arrive on an EVPN device that is not the Ethernet segment's designated forwarder (DF).
 - An IGMP or MLD message to leave a multicast group arrives at a different EVPN device than the EVPN device where the corresponding join message for the group was received.
- Selective multicast forwarding conserves bandwidth usage in the EVPN core and reduces the load on egress EVPN devices that do not have listeners.
- The support of external PIM gateways enables the exchange of multicast traffic between sources and listeners in an EVPN-VXLAN network and sources and listeners in an external PIM domain. Without this support, the sources and listeners in these two domains would not be able to communicate.

Supported IGMP or MLD Versions and Group Membership Report Modes

Table 4 on page 106 outlines the supported IGMP versions and the membership report modes supported for each version.

Table 4: Supported IGMP and MLD Versions and Group Membership Report Modes

IGMP Versions	Any-Source Multicast (ASM) (*,G) Only	Source-Specific Multicast (SSM) (S,G) Only	ASM (*,G) + SSM (S,G)
IGMPv2	Yes (default)	No	No
IGMPv3	Yes (default)	Yes (if configured)	No
MLDv1	Yes (default)	No	No
MLDv2	Yes (default)	Yes (if configured)	No

To explicitly configure EVPN devices to process only SSM (S,G) membership reports for IGMPv3 or MLDv2, set the `evpn-ssm-reports-only` configuration option at the `[edit protocols igmp-snooping vlan vlan-name]` hierarchy level.

You can enable SSM-only processing for one or more VLANs in an EVPN routing instance (EVI). When enabling this option for a routing instance of type `virtual switch`, the behavior applies to all VLANs in the virtual switch instance. When you enable this option, the device doesn't process ASM reports and drops them.

If you don't configure the `evpn-ssm-reports-only` option, by default, EVPN devices process IGMPv2, IGMPv3, MLDv1, or MLDv2 ASM reports and drop IGMPv3 or MLDv2 SSM reports.

Summary of Multicast Traffic Forwarding and Routing Use Cases

[Table 5 on page 107](#) provides a summary of the multicast traffic forwarding and routing use cases that we support in EVPN-VXLAN networks and our recommendation for when you should apply a use case to your EVPN-VXLAN network.

Table 5: Supported Multicast Traffic Forwarding and Routing Use Cases and Recommended Usage

Use Case Number	Use Case Name	Summary	Recommended Usage
1	Intra-VLAN multicast traffic forwarding	Forwarding of multicast traffic to hosts within the same VLAN.	We recommend implementing this basic use case in all EVPN-VXLAN networks.
2	Inter-VLAN multicast routing and forwarding—IRB interfaces with PIM	IRB interfaces using PIM on Layer 3 EVPN devices. These interfaces route multicast traffic between source and receiver VLANs.	We recommend implementing this basic use case in all EVPN-VXLAN networks except when you prefer to use an external multicast router to handle inter-VLAN routing (see use case 5).
3	Inter-VLAN multicast routing and forwarding—PIM gateway with Layer 2 connectivity	A Layer 2 mechanism for a data center, which uses IGMP and PIM, to exchange multicast traffic with an external PIM domain.	We recommend this use case in either EVPN-VXLAN ERB overlays or EVPN-VXLAN CRB overlays.

Table 5: Supported Multicast Traffic Forwarding and Routing Use Cases and Recommended Usage
(Continued)

Use Case Number	Use Case Name	Summary	Recommended Usage
4	Inter-VLAN multicast routing and forwarding—PIM gateway with Layer 3 connectivity	A Layer 3 mechanism for a data center, which uses IGMP (or MLD) and PIM, to exchange multicast traffic with an external PIM domain.	We recommend this use case in EVPN-VXLAN CRB overlays only.
5	Inter-VLAN multicast routing and forwarding—external multicast router	Instead of IRB interfaces on Layer 3 EVPN devices, an external multicast router handles inter-VLAN routing.	We recommend this use case when you prefer to use an external multicast router instead of IRB interfaces on Layer 3 EVPN devices to handle inter-VLAN routing.

For example, in a typical EVPN-VXLAN ERB overlay, you can implement use case 1 for intra-VLAN forwarding and use case 2 for inter-VLAN routing and forwarding. Or, if you want an external multicast router to handle inter-VLAN routing in your EVPN-VXLAN network instead of EVPN devices with IRB interfaces running PIM, you can implement use case 5 instead of use case 2. If there are hosts in an existing external PIM domain that you want hosts in your EVPN-VXLAN network to communicate with, you can also implement use case 3.

When implementing any of the use cases in an EVPN-VXLAN CRB overlay, you can use a mix of spine devices—for example, MX Series routers, EX9200 switches, and QFX10000 switches. However, if you do this, keep in mind that the functionality of all spine devices is determined by the limitations of each spine device. For example, QFX10000 switches support a single routing instance of type `virtual-switch`. Although MX Series routers and EX9200 switches support multiple routing instances of type `evpn` or `virtual-switch`, on each of these devices, you would have to configure a single routing instance of type `virtual-switch` to interoperate with the QFX10000 switches.

Use Case 1: Intra-VLAN Multicast Traffic Forwarding

We recommend this basic use case for all EVPN-VXLAN networks.

This use case supports the forwarding of multicast traffic to hosts within the same VLAN and includes the following key features:

- Hosts that are single-homed to an EVPN device or multihomed to more than one EVPN device in all-active mode.



NOTE: EVPN-VXLAN multicast uses special IGMP and MLD group leave processing to handle multihomed sources and receivers, so we don't support the `immediate-leave` configuration option in the `[edit protocols igmp-snooping]` or `[edit protocols mld-snooping]` hierarchies in EVPN-VXLAN networks.

- Routing instances:
 - The default switch instance if all of the devices in the network support EVPN using the default switch instance.
 - If all devices in the network support EVPN using the `virtual-switch` instance type:
 - (QFX Series switches) A single routing instance of type `virtual-switch`.
 - (MX Series routers, vMX virtual routers, and EX9200 switches) Multiple routing instances of type `evpn` or `virtual-switch`.
 - (Preferred) One or more routing instances of type `mac-vrf` if all devices in the network support EVPN using MAC-VRF instances. (See [MAC-VRF Routing Instance Type Overview](#).)

EVI route target extended community attributes associated with multihomed EVIs. BGP EVPN Type 7 (Join Sync Route) and Type 8 (Leave Synch Route) routes carry these attributes to enable the simultaneous support of multiple EVPN routing instances.

For information about the multicast flags extended community, see ["EVPN Multicast Flags Extended Community" on page 119](#).

- IGMPv2, IGMPv3, MLDv1 or MLDv2. For information about the membership report modes supported for each IGMP or MLD version, see [Table 4 on page 106](#). For information about IGMP or MLD route synchronization between multihomed EVPN devices, see [Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment](#).
- IGMP snooping or MLD snooping. Hosts in a network send IGMP reports (for IPv4 traffic) or MLD reports (for IPv6 traffic) expressing interest in particular multicast groups from multicast sources. EVPN devices with IGMP snooping or MLD snooping enabled listen to the IGMP or MLD reports, and use the snooped information on the access side to establish multicast routes that only forward traffic for a multicast group to interested receivers.

IGMP snooping or MLD snooping supports multicast senders and receivers in the same or different sites. A site can have either receivers only, sources only, or both senders and receivers attached to it.

- Selective multicast forwarding (advertising EVPN Type 6 Selective Multicast Ethernet Tag (SMET) routes for forwarding only to interested receivers). This feature enables EVPN devices to selectively forward multicast traffic to only the devices in the EVPN core that have expressed interest in that multicast group.



NOTE: We support selective multicast forwarding to devices in the EVPN core only in EVPN-VXLAN CRB overlays.

When you enable IGMP snooping or MLD snooping, selective multicast forwarding is enabled by default.

- EVPN devices that do not support IGMP snooping, MLD snooping, and selective multicast forwarding.

Although you can implement this use case in an EVPN single-homed environment, this use case is particularly effective in an EVPN multihomed environment with a high volume of multicast traffic.

All multihomed interfaces must have the same configuration, and all multihomed peer EVPN devices must be in active mode (not standby or passive mode).

An EVPN device that initially receives traffic from a multicast source is known as the *ingress device*. The ingress device handles the forwarding of intra-VLAN multicast traffic as follows:

- With IGMP snooping or MLD snooping enabled (which also enable selective multicast forwarding on supporting devices):
 - As shown in [Figure 9 on page 111](#), the ingress device (leaf 1) selectively forwards the traffic to other EVPN devices with access interfaces where there are interested receivers for the same multicast group.
 - The traffic is then selectively forwarded to egress devices in the EVPN core that have advertised the EVPN Type 6 SMET routes.
- If any EVPN devices do not support IGMP snooping or MLD snooping, or the ability to originate EVPN Type 6 SMET routes, the ingress device floods multicast traffic to these devices.
- If a host is multihomed to more than one EVPN device, the EVPN devices exchange EVPN Type 7 and Type 8 routes as shown in [Figure 9 on page 111](#). This exchange synchronizes IGMP or MLD membership reports received on multihomed interfaces to coordinate status from messages that go to different EVPN devices or in case one of the EVPN devices fails.

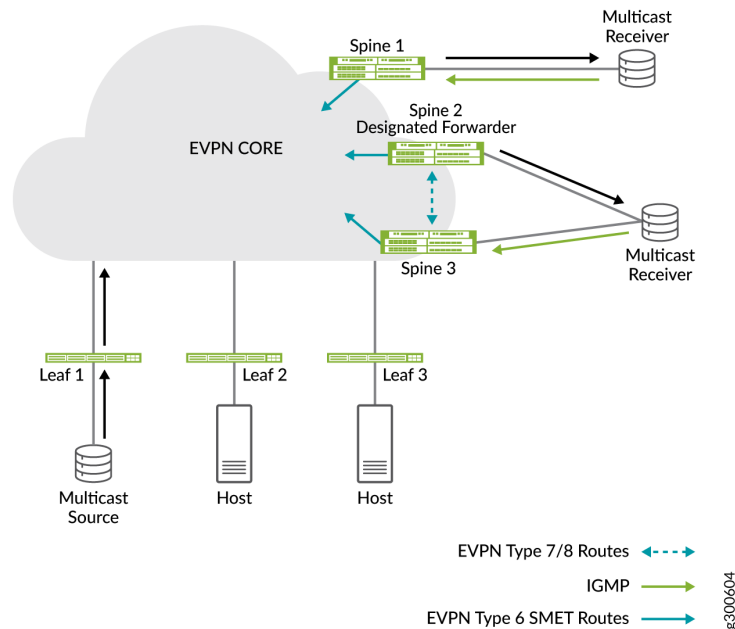


NOTE: The EVPN Type 7 and Type 8 routes carry EVI route extended community attributes to ensure the right EVPN instance gets the IGMP state information on devices with multiple routing instances. QFX Series switches support IGMP snooping only in the default EVPN routing instance (`default-switch`). In Junos OS releases before 17.4R2, 17.3R3, or 18.1R1, these switches did not include EVI route extended community attributes in Type 7 and Type 8 routes, so they don't properly synchronize the IGMP state if you also have other routing instances configured. Starting in Junos OS

releases 17.4R2, 17.3R3, and 18.1R1, QFX10000 switches include the EVI route extended community attributes that identify the target routing instance, and can synchronize IGMP state if IGMP snooping is enabled in the default EVPN routing instance when other routing instances are configured.

In releases that support MLD and MLD snooping in EVPN-VXLAN fabrics with multihoming, the same behavior applies to synchronizing the MLD state.

Figure 9: Intra-VLAN Multicast Traffic Flow with IGMP Snooping and Selective Multicast Forwarding



If you have configured IRB interfaces with PIM on one or more of the Layer 3 devices in your EVPN-VXLAN network (use case 2), note that the ingress device forwards the multicast traffic to the Layer 3 devices. The ingress device takes this action to register itself with the Layer 3 device that acts as the PIM rendezvous point (RP).

Use Case 2: Inter-VLAN Multicast Routing and Forwarding—IRB Interfaces with PIM

We recommend this basic use case for all EVPN-VXLAN networks except when you prefer to use an external multicast router to handle inter-VLAN routing (see Use Case 5: Inter-VLAN Multicast Routing and Forwarding—External Multicast Router).

For this use case, IRB interfaces using Protocol Independent Multicast (PIM) route multicast traffic between source and receiver VLANs. The EVPN devices on which the IRB interfaces reside then forward the routed traffic using these key features:

- Inclusive multicast forwarding with ingress replication
- IGMP snooping or MLD snooping (if supported)
- Selective multicast forwarding

The default behavior of inclusive multicast forwarding is to replicate multicast traffic and flood the traffic to all devices. For this use case, however, we support inclusive multicast forwarding coupled with IGMP snooping (or MLD snooping) and selective multicast forwarding. As a result, the multicast traffic is replicated but selectively forwarded to access interfaces and devices in the EVPN core that have interested receivers.

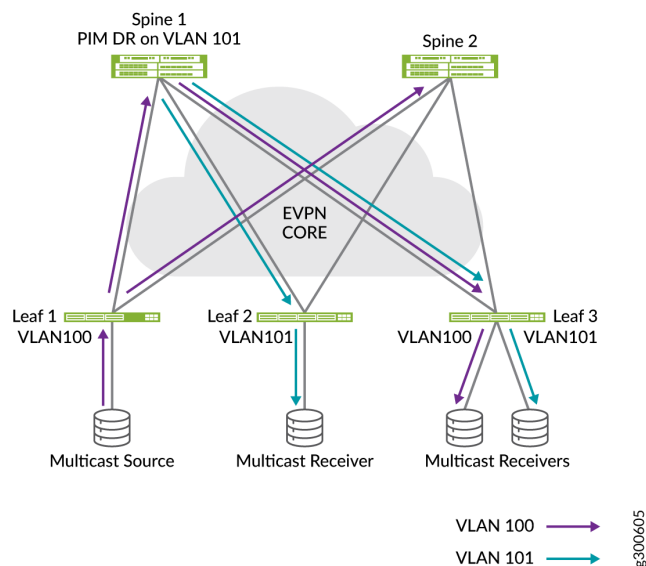
For information about the EVPN multicast flags extended community, which Juniper Networks devices that support EVPN and IGMP snooping (or MLD snooping) include in EVPN Type 3 (Inclusive Multicast Ethernet Tag) routes, see ["EVPN Multicast Flags Extended Community" on page 119](#).

In an EVPN-VXLAN CRB overlay, you can configure the spine devices so that some of them perform inter-VLAN routing and forwarding of multicast traffic and some do not. At a minimum, we recommend that you configure two spine devices to perform inter-VLAN routing and forwarding.

When there are multiple devices that can perform the inter-VLAN routing and forwarding of multicast traffic, one device is elected as the designated router (DR) for each VLAN.

In the sample EVPN-VXLAN CRB overlay shown in [Figure 10 on page 112](#), assume that multicast traffic needs to be routed from source VLAN 100 to receiver VLAN 101. Receiver VLAN 101 is configured on spine 1, which is designated as the DR for that VLAN.

Figure 10: Inter-VLAN Multicast Traffic Flow with IRB Interface and PIM



After the inter-VLAN routing occurs, the EVPN device forwards the routed traffic to:

- Access interfaces that have multicast listeners (IGMP snooping or MLD snooping).
- Egress devices in the EVPN core that have sent EVPN Type 6 SMET routes for the multicast group members in receiver VLAN 2 (selective multicast forwarding).

To understand how IGMP snooping (or MLD snooping) and selective multicast forwarding reduce the impact of the replicating and flooding behavior of inclusive multicast forwarding, assume that an EVPN-VXLAN CRB overlay includes the following elements:

- 100 IRB interfaces using PIM starting with irb.1 and going up to irb.100
- 100 VLANs
- 20 EVPN devices

For the sample EVPN-VXLAN CRB overlay, m represents the number of VLANs, and n represents the number of EVPN devices. Assuming that IGMP snooping (or MLD snooping) and selective multicast forwarding are disabled, when multicast traffic arrives on irb.1, the EVPN device replicates the traffic $m * n$ times or $100 * 20$ times, which equals a rate of 20,000 packets. If the incoming traffic rate for a particular multicast group is 100 packets per second (pps), the EVPN device would have to replicate 200,000 pps for that multicast group.

If IGMP snooping (or MLD snooping) and selective multicast forwarding are enabled in the sample EVPN-VXLAN CRB overlay, assume that there are interested receivers for a particular multicast group on only 4 VLANs and 3 EVPN devices. In this case, the EVPN device replicates the traffic at a rate of $100 * m * n$ times ($100 * 4 * 3$), which equals 1200 pps. Note the significant reduction in the replication rate and the amount of traffic that must be forwarded.

When implementing this use case, keep in mind that there are important differences for EVPN-VXLAN CRB overlays and EVPN-VXLAN ERB overlays. [Table 6 on page 114](#) outlines these differences

Table 6: Use Case 2: Important Differences for EVPN-VXLAN Edge-routed and Centrally-routed Bridging Overlays

EVPN VXLAN IP Fabric Architectures	Support Mix of Juniper Networks Devices?	All EVPN Devices Required to Host All VLANs In EVPN-VXLAN Network?	All EVPN Devices Required to Host All VLANs that Include Multicast Listeners?	Required PIM Configuration
EVPN-VXLAN ERB overlay	No. We support only QFX10000 switches for all EVPN devices.	Yes	Yes	Configure PIM distributed designated router (DDR) functionality on the IRB interfaces of the EVPN devices.
EVPN-VXLAN CRB overlay	<p>Yes.</p> <p>Spine devices: We support mix of MX Series routers, EX9200 switches, and QFX10000 switches.</p> <p>Leaf devices: We support mix of MX Series routers and QFX5110 switches.</p> <p>NOTE: If you deploy a mix of spine devices, keep in mind that the functionality of all spine devices is determined by the limitations of each spine device.</p>	No	No. However, you must configure all VLANs that include multicast listeners on each spine device that performs inter-VLAN routing. You don't need to configure all VLANs that include multicast listeners on each leaf device.	Do not configure DDR functionality on the IRB interfaces of the spine devices. By not enabling DDR on an IRB interface, PIM remains in a default mode on the interface, which means that the interface acts the designated router for the VLANs.

In addition to the differences described in [Table 6 on page 114](#), a hair pinning issue exists with an EVPN-VXLAN CRB overlay. Multicast traffic typically flows from a source host to a leaf device to a spine device, which handles the inter-VLAN routing. The spine device then replicates and forwards the traffic to VLANs and EVPN devices with multicast listeners. When forwarding the traffic in this type of EVPN-VXLAN overlay, be aware that the spine device returns the traffic to the leaf device from which the traffic originated (hair-pinning). This issue is inherent with the design of the EVPN-VXLAN CRB overlay.

When designing your EVPN-VXLAN overlay, keep this issue in mind especially if you expect the volume of multicast traffic in your overlay to be high and the replication rate of traffic ($m * n$ times) to be large.

Use Case 3: Inter-VLAN Multicast Routing and Forwarding—PIM Gateway with Layer 2 Connectivity

We recommend the PIM gateway with Layer 2 connectivity use case for both EVPN-VXLAN ERB overlays and EVPN-VXLAN CRB overlays.

For this use case, we assume the following:

- You have deployed a EVPN-VXLAN network to support a data center.
- In this network, you have already set up:
 - Intra-VLAN multicast traffic forwarding as described in use case 1.
 - Inter-VLAN multicast traffic routing and forwarding as described in use case 2.
- There are multicast sources and receivers within the data center that you want to communicate with multicast sources and receivers in an external PIM domain.



NOTE: We support this use case with both EVPN-VXLAN ERB overlays and EVPN-VXLAN CRB overlays.

The use case provides a mechanism for the data center, which uses IGMP (or MLD) and PIM, to exchange multicast traffic with the external PIM domain. Using a Layer 2 multicast VLAN (MVLAN) and associated IRB interfaces on the EVPN devices in the data center to connect to the PIM domain, you can enable the forwarding of multicast traffic from:

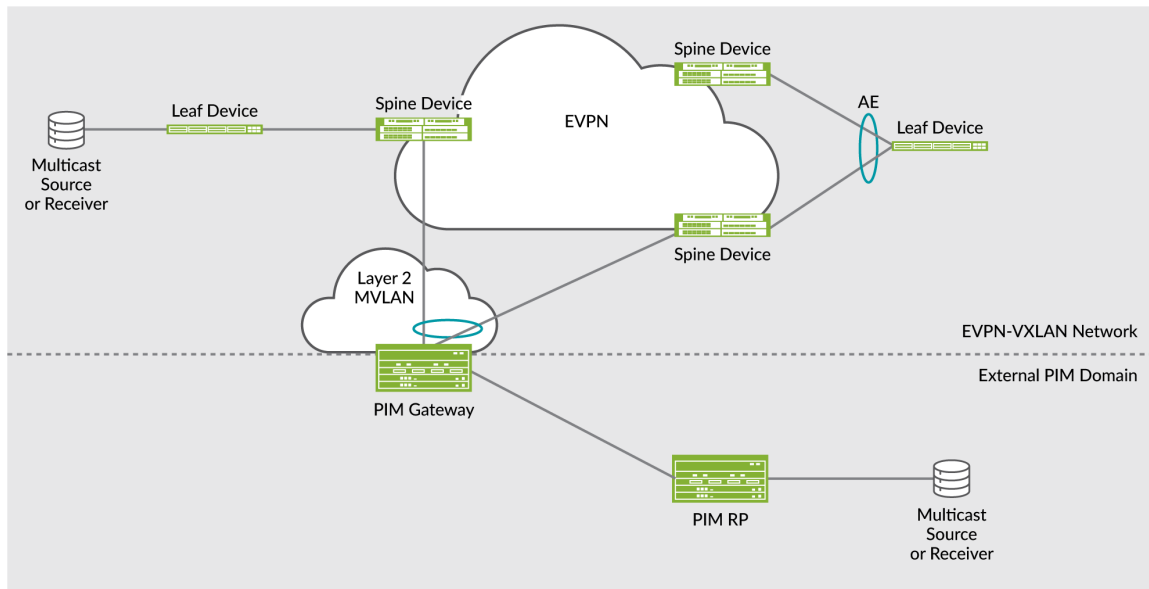
- An external multicast source to internal multicast destinations
- An internal multicast source to external multicast destinations



NOTE: In this section, *external* refers to components in the PIM domain. *Internal* refers to components in your EVPN-VXLAN network that supports a data center.

Figure 11 on page 116 shows the required key components for this use case in a sample EVPN-VXLAN CRB overlay.

Figure 11: Use Case 3: PIM Gateway with Layer 2 Connectivity—Key Components



- Components in the PIM domain:
 - A PIM gateway that acts as an interface between an existing PIM domain and the EVPN-VXLAN network. The PIM gateway is a Juniper Networks or third-party Layer 3 device on which PIM and a routing protocol such as OSPF are configured. The PIM gateway does not run EVPN. You can connect the PIM gateway to one, some, or all EVPN devices.
 - A PIM rendezvous point (RP) is a Juniper Networks or third-party Layer 3 device on which PIM and a routing protocol such as OSPF are configured. You must also configure the PIM RP to translate PIM join or prune messages into corresponding IGMP (or MLD) report or leave messages then forward the reports and messages to the PIM gateway.
- Components in the EVPN-VXLAN network:



NOTE: These components are in addition to the components already configured for use cases 1 and 2.

- EVPN devices. For redundancy, we recommend multihoming the EVPN devices to the PIM gateway through an aggregated Ethernet interface on which you configure an Ethernet segment identifier (ESI). On each EVPN device, you must also configure the following for this use case:
 - A Layer 2 multicast VLAN (MVLAN). The MVLAN is a VLAN that is used to connect the PIM gateway. In the MVLAN, PIM is enabled.

- An MVLAN IRB interface on which you configure PIM, IGMP snooping (or MLD snooping), and a routing protocol such as OSPF. To reach the PIM gateway, the EVPN device forwards multicast traffic out of this interface.
- To enable the EVPN devices to forward multicast traffic to the external PIM domain, configure:
 - PIM-to-IGMP translation:

For EVPN-VXLAN **ERB overlays**, configure PIM-to-IGMP translation by including `pim-to-igmp-proxy upstream-interface irb-interface-name` configuration statements at the [edit routing-options multicast] hierarchy level. Specify the MVLAN IRB interface for the IRB interface parameter. You also must set IGMP passive mode using `igmp interface irb-interface-name passive` configuration statements at the [edit protocols] hierarchy level on the upstream interfaces where you set `pim-to-igmp-proxy`.

For EVPN-VXLAN **CRB overlays**, you do not need to include the `pim-to-igmp-proxy upstream-interface irb-interface-name` or `pim-to-ml-d-proxy upstream-interface irb-interface-name` configuration statements. In this type of overlay, the PIM protocol handles the routing of multicast traffic from the PIM domain to the EVPN-VXLAN network and vice versa.

- Multicast router interface:

Configure the multicast router interface by including the `multicast-router-interface` configuration statement at the [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols (igmp-snooping | mld-snooping) interface *interface-name*] hierarchy level. For the interface name, specify the MVLAN IRB interface.
- PIM passive mode. For EVPN-VXLAN ERB overlays only, you must ensure that the PIM gateway views the data center as only a Layer 2 multicast domain. To do so, include the `passive` configuration statement at the [edit protocols pim] hierarchy level.

Use Case 4: Inter-VLAN Multicast Routing and Forwarding—PIM Gateway with Layer 3 Connectivity

We recommend the PIM gateway with Layer 3 connectivity use case for EVPN-VXLAN CRB overlays only.

For this use case, we assume the following:

- You have deployed a EVPN-VXLAN network to support a data center.
- In this network, you have already set up:
 - Intra-VLAN multicast traffic forwarding as described in use case 1.
 - Inter-VLAN multicast traffic routing and forwarding as described in use case 2.

- There are multicast sources and receivers within the data center that you want to communicate with multicast sources and receivers in an external PIM domain.



NOTE: We recommend the PIM gateway with Layer 3 connectivity use case for EVPN-VXLAN CRB overlays only.

This use case provides a mechanism for the data center, which uses IGMP (or MLD) and PIM, to exchange multicast traffic with the external PIM domain. Using Layer 3 interfaces on the EVPN devices in the data center to connect to the PIM domain, you can enable the forwarding of multicast traffic from:

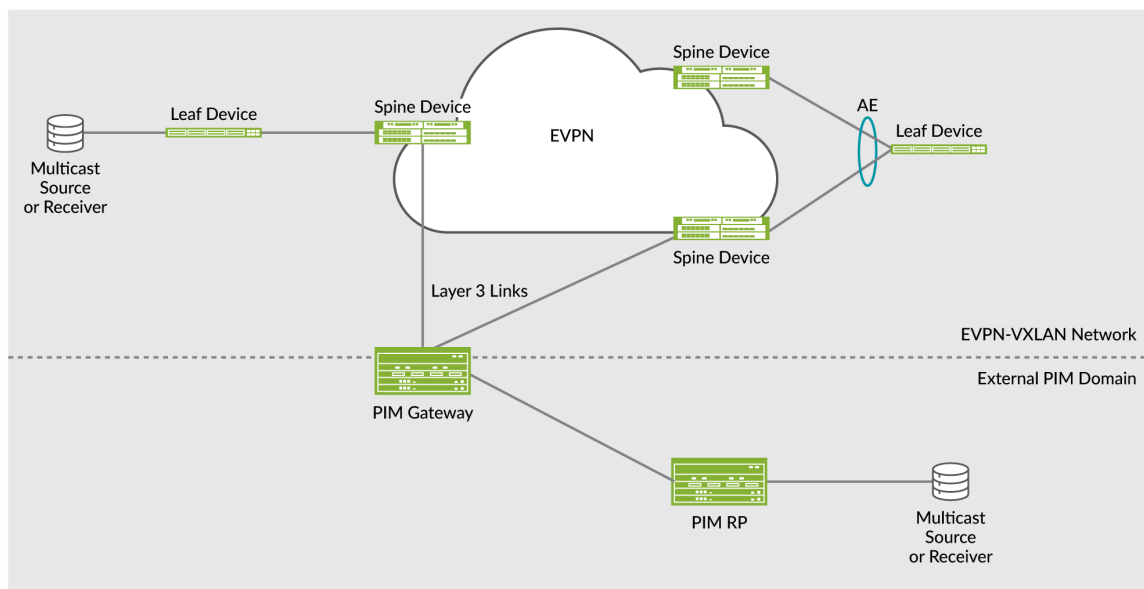
- An external multicast source to internal multicast destinations
- An internal multicast source to external multicast destinations



NOTE: In this section, *external* refers to components in the PIM domains. *Internal* refers to components in your EVPN-VXLAN network that supports a data center.

Figure 12 on page 118 shows the required key components for this use case in a sample EVPN-VXLAN CRB overlay.

Figure 12: Use Case 4: PIM Gateway with Layer 3 Connectivity—Key Components



- Components in the PIM domain:

- A PIM gateway that acts as an interface between an existing PIM domain and the EVPN-VXLAN network. The PIM gateway is a Juniper Networks or third-party Layer 3 device on which PIM and a routing protocol such as OSPF are configured. The PIM gateway does not run EVPN. You can connect the PIM gateway to one, some, or all EVPN devices.
- A PIM rendezvous point (RP) is a Juniper Networks or third-party Layer 3 device on which PIM and a routing protocol such as OSPF are configured. You must also configure the PIM RP to translate PIM join or prune messages into corresponding IGMP or MLD report or leave messages then forward the reports and messages to the PIM gateway.
- Components in the EVPN-VXLAN network:



NOTE: These components are in addition to the components already configured for use cases 1 and 2.

- EVPN devices. You can connect one, some, or all EVPN devices to a PIM gateway. You must make each connection through a Layer 3 interface on which PIM is configured. Other than the Layer 3 interface with PIM, this use case does not require additional configuration on the EVPN devices.

Use Case 5: Inter-VLAN Multicast Routing and Forwarding—External Multicast Router

Starting with Junos OS Release 17.3R1, you can configure an EVPN device to perform inter-VLAN forwarding of multicast traffic without having to configure IRB interfaces on the EVPN device. In such a scenario, an external multicast router is used to send IGMP or MLD queries to solicit reports and to forward VLAN traffic through a Layer 3 multicast protocol such as PIM. IRB interfaces are not supported with the use of an external multicast router.

For this use case, you must include the `igmp-snooping proxy` or `mld-snooping proxy` configuration statements at the `[edit routing-instances routing-instance-name protocols vlan vlan-name]` hierarchy level.

EVPN Multicast Flags Extended Community

Juniper Networks devices that support EVPN-VXLAN and IGMP snooping also support the EVPN multicast flags extended community. When you have enabled IGMP snooping on one of these devices, the device adds the community to EVPN Type 3 (Inclusive Multicast Ethernet Tag) routes.

The absence of this community in an EVPN Type 3 route can indicate the following about the device that advertises the route:

- The device does not support IGMP snooping.
- The device does not have IGMP snooping enabled on it.
- The device is running a Junos OS software release that doesn't support the community.

- The device does not support the advertising of EVPN Type 6 SMET routes.
- The device has IGMP snooping and a Layer 3 interface with PIM enabled on it. Although the Layer 3 interface with PIM performs snooping on the access side and selective multicast forwarding on the EVPN core, the device needs to attract all traffic to perform source registration to the PIM RP and inter-VLAN routing.


The behavior described above also applies to devices that support EVPN-VXLAN with MLD and MLD snooping.

Figure 13 on page 120 shows the EVPN multicast flag extended community, which has the following characteristics:

- The community is encoded as an 8-bit value.
- The Type field has a value of 6.
- The IGMP Proxy Support flag is set to 1, which means that the device supports IGMP proxy.

The same applies to the MLD Proxy Support flag; if that flag is set to 1, the device supports MLD proxy. Either or both flags might be set.

Figure 13: EVPN Multicast Flag Extended Community



Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
22.3R1	Starting with Junos OS Release 22.3R1, we support IGMP snooping on EX4400 switches in EVPN-VXLAN networks.
22.3R1	Starting with Junos release 22.3R1, on some platforms we support IPv4 and IPv6 multicast traffic in an IPv6 EVPN-VXLAN overlay with IPv6 underlay peering.

21.2R1	Starting with Junos OS Release 21.2R1, we support optimized inter-subnet multicast (OISM) routing and forwarding with IGMP snooping on QFX5110, QFX5120, and QFX10002 switches in EVPN-VXLAN ERB overlay networks.
20.4R1	Starting with Junos OS Release 20.4R1, QFX5120-48YM switches support IGMP snooping as leaf devices in a multihomed EVPN-VXLAN CRB overlay fabric.
20.4R1	Starting in Junos OS Release 20.4R1, in EVPN-VXLAN CRB overlay fabrics, QFX5110, QFX5120 and the QFX10000 line of switches support IGMPv3 with IGMP snooping for IPv4 multicast traffic, and MLDv1 and MLDv2 with MLD snooping for IPv6 multicast traffic. Supported switches in the QFX5000 line only support multicast forwarding as leaf devices in CRB fabrics, and don't support multicast routing. The spine devices handle Layer 3 multicast routing. You can configure these switches to process IGMPv3 and MLDv2 source-specific multicast (SSM) reports, but these devices can't process both SSM reports and any-source multicast (ASM) reports at the same time.
20.2R1	Starting with Junos OS Release 20.2R1, QFX5120-48T switches support IGMP snooping as leaf devices in a multihomed EVPN-VXLAN CRB overlay fabric.
19.3R1	Starting with Junos OS Release 19.3R1, EX9200 switches, MX Series routers, and vMX virtual routers support IGMPv2, IGMPv3, IGMP snooping, selective multicast forwarding, external PIM gateways, and external multicast routers with an EVPN-VXLAN CRB overlay.
19.1R1	Starting with Junos OS Release 19.1R1, QFX5120-32C switches support IGMP snooping as leaf devices in an EVPN-VXLAN CRB overlay fabric.
18.4R2	Starting with Junos OS Release 18.4R2 (but not Junos OS Releases 19.1R1 and 19.2R1), QFX5120-48Y switches support IGMP snooping as leaf devices in an EVPN-VXLAN CRB overlay fabric.
18.4R2	Starting in Junos OS Releases 18.4R2 and 19.1R2, selective multicast forwarding is enabled by default on QFX5110 and QFX5120 switches when you configure IGMP snooping in EVPN-VXLAN networks, further constraining multicast traffic flooding. With IGMP snooping and selective multicast forwarding, these switches send the multicast traffic only to interested receivers in both the EVPN core and on the access side for multicast traffic coming in either from an access interface or an EVPN network interface.
18.1R1	Starting with Junos OS Release 18.1R1, QFX5110 switches support IGMP snooping as leaf devices in an EVPN-VXLAN CRB overlay (EVPN-VXLAN topology with a two-layer IP fabric) for forwarding multicast traffic within VLANs. You can't configure IRB interfaces on a VXLAN with IGMP snooping for forwarding multicast traffic between VLANs. You can only configure and use IRB interfaces for unicast traffic.

17.3R1	Starting with Junos OS Release 17.3R1, QFX10000 switches support exchanging traffic between multicast sources and receivers in an EVPN-VXLAN ERB overlay using IGMP. These switches also support multicast traffic flow from sources and to receivers in an external Protocol Independent Multicast (PIM) domain. A Layer 2 multicast VLAN (M-VLAN) and associated IRB interfaces enable the exchange of multicast traffic between these two domains.
17.3R1	Starting with Junos OS Release 17.3R1, you can configure an EVPN device to perform multicast traffic inter-VLAN forwarding without needing to configure IRB interfaces on the EVPN device.
17.2R1	Starting with Junos OS Release 17.2R1, QFX10000 switches support IGMP snooping in an EVPN-VXLAN ERB overlay.

RELATED DOCUMENTATION

[distributed-dr](#)

[igmp-snooping](#)

[mld-snooping](#)

[multicast-router-interface](#)

Example: Preserving Bandwidth with IGMP Snooping in an EVPN-VXLAN Environment

Configuring IGMP Snooping on Switches

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a device. With IGMP snooping enabled, the device monitors IGMP traffic on the network and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. The device conserves bandwidth by sending multicast traffic only to interfaces connected to devices that want to receive the traffic, instead of flooding the traffic to all the downstream interfaces in a VLAN.



NOTE: You cannot configure IGMP snooping on a secondary (private) VLAN (PVLAN). However, starting in Junos OS Release 18.3R1 on EX4300 switches and EX4300 Virtual Chassis, and Junos OS Release 19.2R1 on EX4300 multigigabit switches, you can configure the `vlan` statement at the `[edit protocols igmp-snooping]` hierarchy level with a primary VLAN, which implicitly enables IGMP snooping on its secondary VLANs and avoids flooding multicast traffic on PVLANS. See "[IGMP Snooping on Private VLANs \(PVLANS\)](#)" on page 96 for details.



NOTE: Starting in Junos OS Releases 14.1X53 and 15.2, QFabric Systems support the `igmp-querier` statement to configure a Node device as an IGMP querier.

The default factory configuration on legacy EX Series switches enables IGMP snooping by default on all VLANs. In this case, you don't need any other configuration for IGMP snooping to work. However, if you want IGMP snooping enabled only on some VLANs, you can either disable the feature on all VLANs and then enable it selectively on the desired VLANs, or simply disable the feature selectively on those where you do not want IGMP snooping. You can also customize other available IGMP snooping options.



TIP: When you configure IGMP snooping using the `vlan all` statement (where supported), any VLAN that is not individually configured for IGMP snooping inherits the `vlan all` configuration. Any VLAN that is individually configured for IGMP snooping, on the other hand, does not inherit the `vlan all` configuration. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration. For example, in the following configuration:

```
protocols {
  igmp-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee-vlan {
      interface ge-0/0/8.0 {
        static {
          group 233.252.0.1;
        }
      }
    }
  }
}
```

all VLANs except `employee-vlan` have a robust count of 8. Because you individually configured `employee-vlan`, its robust count value is not determined by the value set under `vlan all`. Instead, its `robust-count` value is 2, the default value.

On switches without IGMP snooping enabled in the default factory configuration, you must explicitly enable IGMP snooping and configure any other of the available IGMP snooping options you want on a VLAN.

Use the following configuration steps as needed for your network to enable IGMP snooping on all VLANs (where supported), enable or disable IGMP snooping selectively on a VLAN, and configure available IGMP snooping options:

1. To enable IGMP snooping on all VLANs (where supported, such as on some EX Series switches):

```
[edit protocols]
user@switch# set igmp-snooping vlan all
```



NOTE: The default factory configuration on legacy EX Series switches has IGMP snooping enabled on all VLANs.

Or disable IGMP snooping on all VLANs (where supported, such as on some EX Series switches):

```
[edit protocols]
user@switch# set igmp-snooping vlan all disable
```

2. To enable IGMP snooping on a specified VLAN, for example, on a VLAN named employee-vlan:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

3. To configure the switch to immediately remove group memberships from interfaces on a VLAN when it receives a leave message through that VLAN, so it doesn't forward any membership queries for the multicast group to the VLAN (IGMPv2 only):

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name immediate-leave
```

4. To configure an interface to statically belong to a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name static group group-address
```

5. To configure an interface to forward IGMP queries it receives from multicast routers:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name multicast-router-
interface
```

6. To change the default number of timeout intervals the device waits before timing out and removing a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count number
```

7. If you want a device to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

Or on QFabric Systems only, if you want a QFabric Node device to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

The switch sends IGMP queries with the configured source address. To ensure this switch is always the IGMP querier on the network, make sure the source address is greater (a higher number) than the IP addresses for any other multicast routers on the same local network.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
14.1X53	Starting in Junos OS Releases 14.1X53 and 15.2, QFabric Systems support the <code>igmp-querier</code> statement to configure a Node device as an IGMP querier.

RELATED DOCUMENTATION

[IGMP Snooping Overview](#) | 96

[Example: Configuring IGMP Snooping on Switches | 126](#)

[Monitoring IGMP Snooping | 139](#)

Example: Configuring IGMP Snooping on Switches

IN THIS SECTION

- [Requirements | 126](#)
- [Overview and Topology | 126](#)
- [Configuration | 127](#)

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a device. With IGMP snooping enabled, the device monitors IGMP traffic on the network and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. The device conserves bandwidth by sending multicast traffic only to interfaces connected to devices that want to receive the traffic, instead of flooding the traffic to all the downstream interfaces in a VLAN.

This example describes how to configure IGMP snooping:

Requirements

This example requires Junos OS Release 11.1 or later on a QFX Series product.

Before you configure IGMP snooping, be sure you have:

- Configured the `employee-vlan` VLAN
- Assigned interfaces `ge-0/0/1`, `ge-0/0/2`, `ge-0/0/3`, and `ge-0/0/4` to `employee-vlan`

Overview and Topology

IN THIS SECTION

- [Topology | 127](#)

In this example you configure an interface to receive multicast traffic from a source and configure some multicast-related behavior for downstream interfaces. The example assumes that IGMP snooping was previously disabled for the VLAN.

By enabling IGMP snooping on a switch, interfaces can be either host-only interfaces or multicast-router interfaces. If the interfaces are not configured explicitly, through IGMP snooping the switch learns which interfaces are host-only and which ones are multicast-router interfaces.

Topology

[Table 7 on page 127](#) shows the components of the topology for this example.

Table 7: Components of the IGMP Snooping Topology

Components	Settings
VLAN name	employee-vlan, tag 20
Interfaces in employee-vlan	ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/4
Multicast IP address for employee-vlan	225.100.100.100

Configuration

IN THIS SECTION

- [Procedure | 128](#)

To configure basic IGMP snooping on a switch:

Procedure

CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into a terminal window:

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN `employee-vlan`:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

2. Optionally, you can configure an interface to belong to a multicast group. For example, for testing IGMP snooping with Layer 2 multicast forwarding, you might assign an interface to a static multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group
225.100.100.100
```

(See [static \(IGMP Snooping\)](#) for more on how static groups work at Layer 2.)

3. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-
interface
```

4. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

Results

Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
vlan employee-vlan {
  robust-count 4;
}
interface ge-0/0/2 {
  multicast-router-interface;
}
interface ge-0/0/3 {
  static {
    group 255.100.100.100;
  }
}
}
```

RELATED DOCUMENTATION

[IGMP Snooping Overview | 96](#)

[Configuring IGMP Snooping on Switches | 122](#)

[Changing the IGMP Snooping Group Timeout Value on Switches | 138](#)

[Monitoring IGMP Snooping | 139](#)

Example: Configuring IGMP Snooping on EX Series Switches

IN THIS SECTION

- [Requirements | 130](#)
- [Overview and Topology | 130](#)
- [Configuration | 132](#)
- [Verifying IGMP Snooping Operation | 133](#)

You can enable IGMP snooping on a VLAN to constrain the flooding of IPv4 multicast traffic on a VLAN. When IGMP snooping is enabled, a switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure IGMP snooping:

Requirements

This example uses the following software and hardware components:

- One EX4300 Series switch
- Junos OS Release 13.2 or later for EX Series switches

Before you configure IGMP snooping, be sure you have:

- Configured the **vlan100** VLAN on the switch
- Assigned interfaces **ge-0/0/0**, **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/12** to vlan100
- Configure **ge-0/0/12** as a trunk interface.

Overview and Topology

IN THIS SECTION

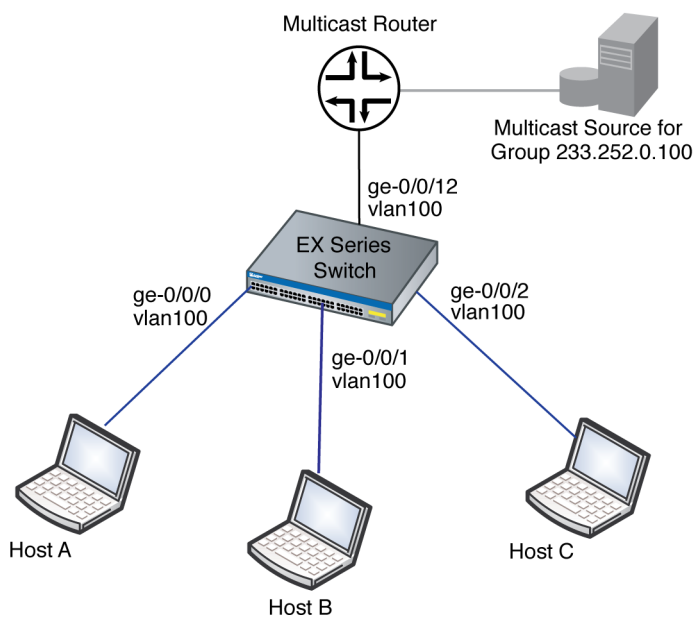
- [Topology | 131](#)

In this example, interfaces **ge-0/0/0**, **ge-0/0/1**, and **ge-0/0/2** on the switch are in **vlan100** and are connected to hosts that are potential multicast receivers. Interface **ge-0/0/12**, a trunk interface also in **vlan100**, is connected to a multicast router. The router acts as the IGMP querier and forwards multicast traffic for group **233.255.0.100** to the switch from a multicast source.

Topology

The example topology is illustrated in [Figure 14 on page 131](#).

Figure 14: Example IGMP Snooping Topology



In this example topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group **233.255.0.100** from one of the hosts—for example, Host B. If IGMP snooping is not enabled on **vlan100**, the switch floods the multicast traffic on all interfaces in **vlan100** (except for interface **ge-0/0/12**). If IGMP snooping is enabled on **vlan100**, the switch monitors the IGMP messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface **ge-0/0/1**.

IGMP snooping is enabled on all VLANs in the default factory configuration. For many implementations, IGMP snooping requires no additional configuration. This example shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast

group has left the group. If immediate leave is not configured, the switch waits until the group-specific queries time out before it stops forwarding traffic.

Immediate leave is supported by IGMP version 2 (IGMPv2) and IGMPv3. With IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface. In IGMPv2, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports to avoid a flood of reports for the same group. This report-suppression feature means that the switch only knows about one interested host at any given time.

- Configure **ge-0/0/12** as a static multicast-router interface. In this topology, **ge-0/0/12** always leads to the multicast router. By statically configuring **ge-0/0/12** as a multicast-router interface, you avoid any delay imposed by the switch having to learn that **ge-0/0/12** is a multicast-router interface.

Configuration

IN THIS SECTION

- [Procedure | 132](#)

To configure IGMP snooping on a switch:

Procedure

CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols igmp-snooping vlan vlan100 immediate-leave
set protocols igmp-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Step-by-Step Procedure

To configure IGMP snooping on vlan100:

1. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 immediate-leave
```

2. Statically configure interface **ge-0/0/12** as a multicast-router interface:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show igmp-snooping
vlan all;
vlan vlan100 {
    immediate-leave;
    interface ge-0/0/12.0 {
        multicast-router-interface;
    }
}
```

Verifying IGMP Snooping Operation

IN THIS SECTION

- [Displaying IGMP Snooping Information for VLAN vlan100 | 134](#)

To verify that IGMP snooping is operating as configured, perform the following task:

Displaying IGMP Snooping Information for VLAN vlan100

Purpose

Verify that IGMP snooping is enabled on **vlan100** and that **ge-0/0/12** is recognized as a multicast-router interface.

Action

Enter the following command:

```
user@switch> show igmp-snooping vlans vlan vlan100 detail
VLAN: vlan100, Tag: 100
  Interface: ge-0/0/12.0, tagged, Groups: 0, Router
```

Meaning

By showing information for **vlan100**, the command output confirms that IGMP snooping is configured on the VLAN. Interface **ge-0/0/12.0** is listed as multicast-router interface, as configured. Because none of the host interfaces are listed, none of the hosts are currently receivers for the multicast group.

RELATED DOCUMENTATION

[Configuring IGMP Snooping on Switches | 122](#)

[Verifying IGMP Snooping on EX Series Switches | 134](#)

[IGMP Snooping Overview | 96](#)

Verifying IGMP Snooping on EX Series Switches

IN THIS SECTION

- [Verifying IGMP Snooping Memberships | 135](#)
- [Viewing IGMP Snooping Statistics | 136](#)

- [Viewing IGMP Snooping Routing Information | 137](#)

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a switch. This topic describes how to verify IGMP snooping operation on the switch.

It covers:

Verifying IGMP Snooping Memberships

IN THIS SECTION

- [Purpose | 135](#)
- [Action | 135](#)
- [Meaning | 136](#)

Purpose

Determine group memberships, multicast-router interfaces, host IGMP versions, and the current values of timeout counters.

Action

Enter the following command:

```
user@switch> show igmp snooping membership detail
VLAN: vlan2 Tag: 2 (Index: 3)
Router interfaces:
  ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
Group: 233.252.0.1
  ge-1/0/17.0 259 Last reporter: 10.0.0.90 Receiver count: 1
  Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
  Include source: 10.2.11.5, 10.2.11.12
```

Meaning

The switch has multicast membership information for one VLAN on the switch, **vlan2**. IGMP snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them. The following information is provided:

- Information on the multicast-router interfaces for the VLAN—in this case, **ge-1/0/0.0**. The multicast-router interface has been learned by IGMP snooping, as indicated by the dynamic value. The timeout value shows how many seconds from now the interface will be removed from the multicast forwarding table if the switch does not receive IGMP queries or Protocol Independent Multicast (PIM) updates on the interface.
- Information about the group memberships for the VLAN:
 - Currently, the VLAN has membership in only one multicast group, **233.252.0.1**.
 - The host or hosts that have reported membership in the group are on interface **ge-1/0/17.0**. The last host that reported membership in the group has address **10.0.0.90**. The number of hosts belonging to the group on the interface is shown in the Receiver count field, which is displayed only when host tracking is enabled if immediate leave is configured on the VLAN.
 - The Uptime field shows that the multicast group has been active on the interface for 19 seconds. The interface group membership will time out in 259 seconds if no hosts respond to membership queries during this interval. The Flags field shows the lowest version of IGMP used by a host that is currently a member of the group, which in this case is IGMP version 3 (IGMPv3).
 - Because the interface has IGMPv3 hosts on it, the source addresses from which the IGMPv3 hosts want to receive group multicast traffic are shown (addresses **10.2.11.5** and **10.2.11.12**). The timeout value for the interface group membership is derived from the largest timeout value for all sources addresses for the group.

Viewing IGMP Snooping Statistics

IN THIS SECTION

- [Purpose | 137](#)
- [Action | 137](#)
- [Meaning | 137](#)

Purpose

Display IGMP snooping statistics, such as number of IGMP queries, reports, and leaves received and how many of these IGMP messages contained errors.

Action

Enter the following command:

```
user@switch> show igmp snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0
```

IGMP Type	Received	Transmitted	Recv Errors
Queries:	74295	0	0
Reports:	18148423	0	16333523
Leaves:	0	0	0
Other:	0	0	0

Meaning

The output shows how many IGMP messages of each type—**Queries**, **Reports**, **Leaves**—the switch received or transmitted on interfaces on which IGMP snooping is enabled. For each message type, it also shows the number of IGMP packets the switch received that had errors—for example, packets that do not conform to the IGMPv1, IGMPv2, or IGMPv3 standards. If the **Recv Errors** count increases, verify that the hosts are compliant with IGMP standards. If the switch is unable to recognize the IGMP message type for a packet, it counts the packet under **Receive unknown**.

Viewing IGMP Snooping Routing Information

IN THIS SECTION

- [Purpose | 138](#)
- [Action | 138](#)
- [Meaning | 138](#)

Purpose

Display the next-hop information maintained in the multicast forwarding table.

Action

Enter the following command:

```
user@switch> show multicast snooping route vlan
```

Meaning

The output shows the next-hop interfaces for a given multicast group on a VLAN.

RELATED DOCUMENTATION

clear igmp snooping membership

[Example: Configuring IGMP Snooping on EX Series Switches | 130](#)

[Configuring IGMP Snooping on Switches | 122](#)

Changing the IGMP Snooping Group Timeout Value on Switches

The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. A switch calculates the timeout value by using the query-interval and query-response-interval values.

When you enable IGMP snooping, the query-interval and query-response-interval values are applied to all VLANs on the switch. The values are:

- query-interval—125 seconds
- query-response-interval—10 seconds

The switch automatically calculates the group timeout value for an IGMP snooping-enabled switch by multiplying the query-interval value by 2 (the default robust-count value) and then adding the query-response-interval value. By default, the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table: $(125 \times 2) + 10 = 260$.

You can modify the group timeout value by changing the `robust-count` value. For example, if you want the system to wait 510 seconds before timing groups out— $(125 \times 4) + 10 = 510$ —enter this command:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

RELATED DOCUMENTATION

[Verifying IGMP Snooping on EX Series Switches | 134](#)

[Example: Configuring IGMP Snooping on Switches | 126](#)

[Configuring IGMP Snooping on Switches | 122](#)

Monitoring IGMP Snooping

IN THIS SECTION

- [Purpose | 139](#)
- [Action | 139](#)
- [Meaning | 140](#)

Purpose

Use the monitoring feature to view status and information about the IGMP snooping configuration.

Action

To display details about IGMP snooping, enter the following operational commands:

- `show igmp snooping interface`—Display information about interfaces enabled with IGMP snooping, including which interfaces are being snooped in a learning domain and the number of groups on each interface.
- `show igmp snooping membership`—Display IGMP snooping membership information, including the multicast group address and the number of active multicast groups.

- `show igmp snooping options`—Display brief or detailed information about IGMP snooping.
- `show igmp snooping statistics`—Display IGMP snooping statistics, including the number of messages sent and received.

The `show igmp snooping interface`, `show igmp snooping membership`, and `show igmp snooping statistics` commands also support the following options:

- `instance instance-name`
- `interface interface-name`
- `qualified-vlan vlan-identifier`
- `vlan vlan-name`

Meaning

[Table 8 on page 140](#) summarizes the IGMP snooping details displayed.

Table 8: Summary of IGMP Snooping Output Fields

Field	Values
IGMP Snooping Monitor	
VLAN	VLAN for which IGMP snooping is enabled.
Interfaces	Interface connected to a multicast router.
Groups	Number of the multicast groups learned by the VLAN.
MRouters	Multicast router.
Receivers	Multicast receiver.
IGMP Route Information	
VLAN	VLAN for which IGMP snooping is enabled.

Table 8: Summary of IGMP Snooping Output Fields (*Continued*)

Field	Values
Next-Hop	Next hop assigned by the switch after performing the route lookup.
Group	Multicast groups learned by the VLAN.

RELATED DOCUMENTATION

[IGMP Snooping Overview | 96](#)

[Example: Configuring IGMP Snooping on Switches | 126](#)

[Configuring IGMP Snooping on Switches | 122](#)

[Changing the IGMP Snooping Group Timeout Value on Switches | 138](#)

Example: Configuring IGMP Snooping

IN THIS SECTION

- [Understanding Multicast Snooping | 142](#)
- [Understanding IGMP Snooping | 142](#)
- [IGMP Snooping Interfaces and Forwarding | 144](#)
- [IGMP Snooping and Proxies | 144](#)
- [Multicast-Router Interfaces and IGMP Snooping Proxy Mode | 145](#)
- [Host-Side Interfaces and IGMP Snooping Proxy Mode | 146](#)
- [IGMP Snooping and Bridge Domains | 146](#)
- [Configuring IGMP Snooping | 147](#)
- [Configuring VLAN-Specific IGMP Snooping Parameters | 148](#)
- [Example: Configuring IGMP Snooping | 149](#)
- [Configuring IGMP Snooping Trace Operations | 158](#)
- [Configuring IGMP or MLD Snooping Version | 161](#)
- [Restrict Flooding of Unknown Multicast Traffic to the Multicast-router Interface in an L2 environment | 163](#)

Understanding Multicast Snooping

Network devices such as routers operate mainly at the packet level, or Layer 3. Other network devices such as bridges or LAN switches operate mainly at the frame level, or Layer 2. Multicasting functions mainly at the packet level, Layer 3, but there is a way to map Layer 3 IP multicast group addresses to Layer 2 MAC multicast group addresses at the frame level.

Routers can handle both Layer 2 and Layer 3 addressing information because the frame and its addresses must be processed to access the encapsulated packet inside. Routers can run Layer 3 multicast protocols such as PIM or IGMP and determine where to forward multicast content or when a host on an interface joins or leaves a group. However, bridges and LAN switches, as Layer 2 devices, are not supposed to have access to the multicast information inside the packets that their frames carry.

How then are bridges and other Layer 2 devices to determine when a device on an interface joins or leaves a multicast tree, or whether a host on an attached LAN wants to receive the content of a particular multicast group?

The answer is for the Layer 2 device to implement multicast snooping. Multicast snooping is a general term and applies to the process of a Layer 2 device “snooping” at the Layer 3 packet content to determine which actions are taken to process or forward a frame. There are more specific forms of snooping, such as IGMP snooping or PIM snooping. In all cases, snooping involves a device configured to function at Layer 2 having access to normally “forbidden” Layer 3 (packet) information. Snooping makes multicasting more efficient in these devices.

SEE ALSO

| [Multicast Overview | 2](#)

Understanding IGMP Snooping

Snooping is a general way for Layer 2 devices, such as Juniper Networks MX Series Ethernet Services Routers, to implement a series of procedures to “snoop” at the Layer 3 packet content to determine which actions are to be taken to process or forward a frame. More specific forms of snooping, such as Internet Group Membership Protocol (IGMP) snooping or Protocol Independent Multicast (PIM) snooping, are used with multicast.

Layer 2 devices (LAN switches or bridges) handle multicast packets and the frames that contain them much in the same way the Layer 3 devices (routers) handle broadcasts. So, a Layer 2 switch processes an arriving frame having a multicast destination media access control (MAC) address by forwarding a copy of the packet (frame) onto each of the other network interfaces of the switch that are in a forwarding state.

However, this approach (sending multicast frames everywhere the device can) is not the most efficient use of network bandwidth, particularly for IPTV applications. IGMP snooping functions by “snooping” at

the IGMP packets received by the switch interfaces and building a multicast database similar to that a multicast router builds in a Layer 3 network. Using this database, the switch can forward multicast traffic only onto downstream interfaces with interested receivers, and this technique allows more efficient use of network bandwidth.

You configure IGMP snooping for each bridge on the router. A bridge instance without qualified learning has just one learning domain. For a bridge instance with qualified learning, snooping will function separately within each learning domain in the bridge. That is, IGMP snooping and multicast forwarding will proceed independently in each learning domain in the bridge.

This discussion focuses on bridge instances without qualified learning (those forming one learning domain on the device). Therefore, all the interfaces mentioned are logical interfaces of the bridge or VPLS instance.

Several related concepts are important when discussing IGMP snooping:

- Bridge or VPLS instance interfaces are either multicast-router interfaces or host-side interfaces.
- IGMP snooping supports proxy mode or without-proxy mode.



NOTE: When integrated routing and bridging (IRB) is used, if the router is an IGMP querier, any leave message received on any Layer 2 interface will cause a group-specific query on all Layer 2 interfaces (as a result of this practice, some corresponding reports might be received on all Layer 2 interfaces). However, if some of the Layer 2 interfaces are also router (Layer 3) interfaces, reports and leaves from other Layer 2 interfaces will not be forwarded on those interfaces.

If an IRB interface is used as an outgoing interface in a multicast forwarding cache entry (as determined by the routing process), then the output interface list is expanded into a subset of the Layer 2 interface in the corresponding bridge. The subset is based on the snooped multicast membership information, according to the multicast forwarding cache entry installed by the snooping process for the bridge.

If no snooping is configured, the IRB output interface list is expanded to all Layer 2 interfaces in the bridge.

The Junos OS does not support IGMP snooping in a VPLS configuration on a virtual switch. This configuration is disallowed in the CLI.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

[Example: Configuring IGMP Snooping | 141](#)

IGMP Snooping Interfaces and Forwarding

IGMP snooping divides the device interfaces into multicast-router interfaces and host-side interfaces. A multicast-router interface is an interface in the direction of a multicasting router. An interface on the bridge is considered a multicast-router interface if it meets at least one of the following criteria:

- It is statically configured as a multicast-router interface in the bridge instance.
- IGMP queries are being received on the interface.

All other interfaces that are not multicast-router interfaces are considered host-side interfaces.

Any multicast traffic received on a bridge interface with IGMP snooping configured will be forwarded according to following rules:

- Any IGMP packet is sent to the Routing Engine for snooping processing.
- Other multicast traffic with destination address 224.0.0/24 is flooded onto all other interfaces of the bridge.
- Other multicast traffic is sent to all the multicast-router interfaces but only to those host-side interfaces that have hosts interested in receiving that multicast group.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

IGMP Snooping and Proxies

Without a proxy arrangement, IGMP snooping does not generate or introduce queries and reports. It will only “snoop” reports received from all of its interfaces (including multicast-router interfaces) to build its state and group (S,G) database.

Without a proxy, IGMP messages are processed as follows:

- Query—All general and group-specific IGMP query messages received on a multicast-router interface are forwarded to all other interfaces (both multicast-router interfaces and host-side interfaces) on the bridge.
- Report—IGMP reports received on any interface of the bridge are forwarded toward other multicast-router interfaces. The receiving interface is added as an interface for that group if a multicast routing entry exists for this group. Also, a group timer is set for the group on that interface. If this timer expires (that is, there was no report for this group during the IGMP group timer period), then the interface is removed as an interface for that group.

- Leave—IGMP leave messages received on any interface of the bridge are forwarded toward other multicast-router interfaces on the bridge. The Leave Group message reduces the time it takes for the multicast router to stop forwarding multicast traffic when there are no longer any members in the host group.

Proxy snooping reduces the number of IGMP reports sent toward an IGMP router.



NOTE: With proxy snooping configured, an IGMP router is not able to perform host tracking.

As proxy for its host-side interfaces, IGMP snooping in proxy mode replies to the queries it receives from an IGMP router on a multicast-router interface. On the host-side interfaces, IGMP snooping in proxy mode behaves as an IGMP router and sends general and group-specific queries on those interfaces.



NOTE: Only group-specific queries are generated by IGMP snooping directly. General queries received from the multicast-router interfaces are flooded to host-side interfaces.

All the queries generated by IGMP snooping are sent using 0.0.0.0 as the source address. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as the source address unless there is a configured source address to use.

Proxy mode functions differently on multicast-router interfaces than it does on host-side interfaces.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

Multicast-Router Interfaces and IGMP Snooping Proxy Mode

On multicast-router interfaces, in response to IGMP queries, IGMP snooping in proxy mode sends reports containing aggregate information on groups learned on all host-side interfaces of the bridge.

Besides replying to queries, IGMP snooping in proxy mode forwards all queries, reports, and leaves received on a multicast-router interface to other multicast-router interfaces. IGMP snooping keeps the membership information learned on this interface but does not send a group-specific query for leave messages received on this interface. It simply times out the groups learned on this interface if there are no reports for the same group within the timer duration.



NOTE: For the hosts on all the multicast-router interfaces, it is the IGMP router, not the IGMP snooping proxy, that generates general and group-specific queries.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

Host-Side Interfaces and IGMP Snooping Proxy Mode

No reports are sent on host-side interfaces by IGMP snooping in proxy mode. IGMP snooping processes reports received on these interfaces and sends group-specific queries onto host-side interfaces when it receives a leave message on the interface. Host-side interfaces do not generate periodic general queries, but forwards or floods general queries received from multicast-router interfaces.

If a group is removed from a host-side interface and this was the last host-side interface for that group, a leave is sent to the multicast-router interfaces. If a group report is received on a host-side interface and this was the first host-side interface for that group, a report is sent to all multicast-router interfaces.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

IGMP Snooping and Bridge Domains

IGMP snooping on a VLAN is only allowed for the legacy **vlan-id all** case. In other cases, there is a specific bridge domain configuration that determines the VLAN-specific configuration for IGMP snooping.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

Configuring IGMP Snooping

To configure Internet Group Management Protocol (IGMP) snooping, include the **igmp-snooping** statement:

```
igmp-snooping {
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    immediate-leave;
    multicast-router-interface;
    static {
      group ip-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  vlan vlan-id {
    immediate-leave;
    interface interface-name {
      group-limit limit;
      host-only-interface;
      immediate-leave;
      multicast-router-interface;
      static {
        group ip-address {
          source ip-address;
        }
      }
    }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
```

```

    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
}
}

```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* protocols]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]

By default, IGMP snooping is not enabled. Statements configured at the VLAN level apply only to that particular VLAN.

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

Configuring VLAN-Specific IGMP Snooping Parameters

All of the IGMP snooping statements configured with the `igmp-snooping` statement, with the exception of the `traceoptions` statement, can be qualified with the same statement at the VLAN level. To configure IGMP snooping parameters at the VLAN level, include the `vlan` statement:

```

vlan vlan-id;
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        multicast-router-interface;
        static {
            group ip-address {
                source ip-address;
            }
        }
    }
    proxy {
        source-address ip-address;
    }
    query-interval seconds;

```

```

query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
}

```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* protocols igmp-snooping]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping]

SEE ALSO

[Multicast Overview | 2](#)

[Example: Configuring IGMP Snooping | 141](#)

Example: Configuring IGMP Snooping

IN THIS SECTION

- [Requirements | 149](#)
- [Overview and Topology | 150](#)
- [Configuration | 154](#)
- [Verification | 158](#)

This example shows how to configure IGMP snooping. IGMP snooping can reduce unnecessary traffic from IP multicast applications.

Requirements

This example uses the following hardware components:

- One MX Series router
- One Layer 3 device functioning as a multicast router

Before you begin:

- Configure the interfaces. See the [Interfaces User Guide for Security Devices](#).

- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure a multicast protocol. This feature works with the following multicast protocols:
 - DVMRP
 - PIM-DM
 - PIM-SM
 - PIM-SSM

Overview and Topology

IN THIS SECTION

- [Topology | 152](#)

IGMP snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the Layer 2 device broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, a Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts.

This example includes the following statements:

- **proxy**—Enables the Layer 2 device to actively filter IGMP packets to reduce load on the multicast router. Joins and leaves heading upstream to the multicast router are filtered so that the multicast router has a single entry for the group, regardless of how many active listeners have joined the group. When a listener leaves a group but other listeners remain in the group, the leave message is filtered because the multicast router does not need this information. The status of the group remains the same from the router's point of view.
- **immediate-leave**—When only one IGMP host is connected, the `immediate-leave` statement enables the multicast router to immediately remove the group membership from the interface and suppress the sending of any group-specific queries for the multicast group.

When you configure this feature on IGMPv2 interfaces, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to a LAN through the same

interface, and one host sends a leave message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the next general multicast listener query from the router.

When IGMP snooping is enabled on a router running IGMP version 3 (IGMPv3) snooping, after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the Junos OS host-tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

- **query-interval**—Enables you to change the number of IGMP messages sent on the subnet by configuring the interval at which the IGMP querier router sends general host-query messages to solicit membership information.

By default, the query interval is 125 seconds. You can configure any value in the range 1 through 1024 seconds.

- **query-last-member-interval**—Enables you to change the amount of time it takes a device to detect the loss of the last member of a group.

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages.

By default, the last-member query interval is 1 second. You can configure any value in the range 0.1 through 0.9 seconds, and then 1-second intervals from 1 through 1024 seconds.

- **query-response-interval**—Configures how long the router waits to receive a response from its host-query messages.

By default, the query response interval is 10 seconds. You can configure any value in the range 1 through 1024 seconds. This interval should be less than the interval set in the `query-interval` statement.

- **robust-count**—Provides fine-tuning to allow for expected packet loss on a subnet. It is basically the number of intervals to wait before timing out a group. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost.

By default, the robust count is 2. You can configure any value in the range 2 through 10 intervals.

- **group-limit**—Configures a limit for the number of multicast groups (or [S,G] channels in IGMPv3) that can join an interface. After this limit is reached, new reports are ignored and all related flows are discarded, not flooded.

By default, there is no limit to the number of groups that can join an interface. You can configure a limit in the range 0 through a 32-bit number.

- **host-only-interface**—Configure an IGMP snooping interface to be an exclusively host-side interface. On a host-side interface, received IGMP queries are dropped.

By default, an interface can face either other multicast routers or hosts.

- **multicast-router-interface**—Configures an IGMP snooping interface to be an exclusively router-facing interface.

By default, an interface can face either other multicast routers or hosts.

- **static**—Configures an IGMP snooping interface with multicast groups statically.

By default, the router learns about multicast groups on the interface dynamically.

Topology

[Figure 15 on page 153](#) shows networks without IGMP snooping. Suppose host A is an IP multicast sender and hosts B and C are multicast receivers. The router forwards IP multicast traffic only to those segments with registered receivers (hosts B and C). However, the Layer 2 devices flood the traffic to all hosts on all interfaces.

Figure 15: Networks Without IGMP Snooping Configured


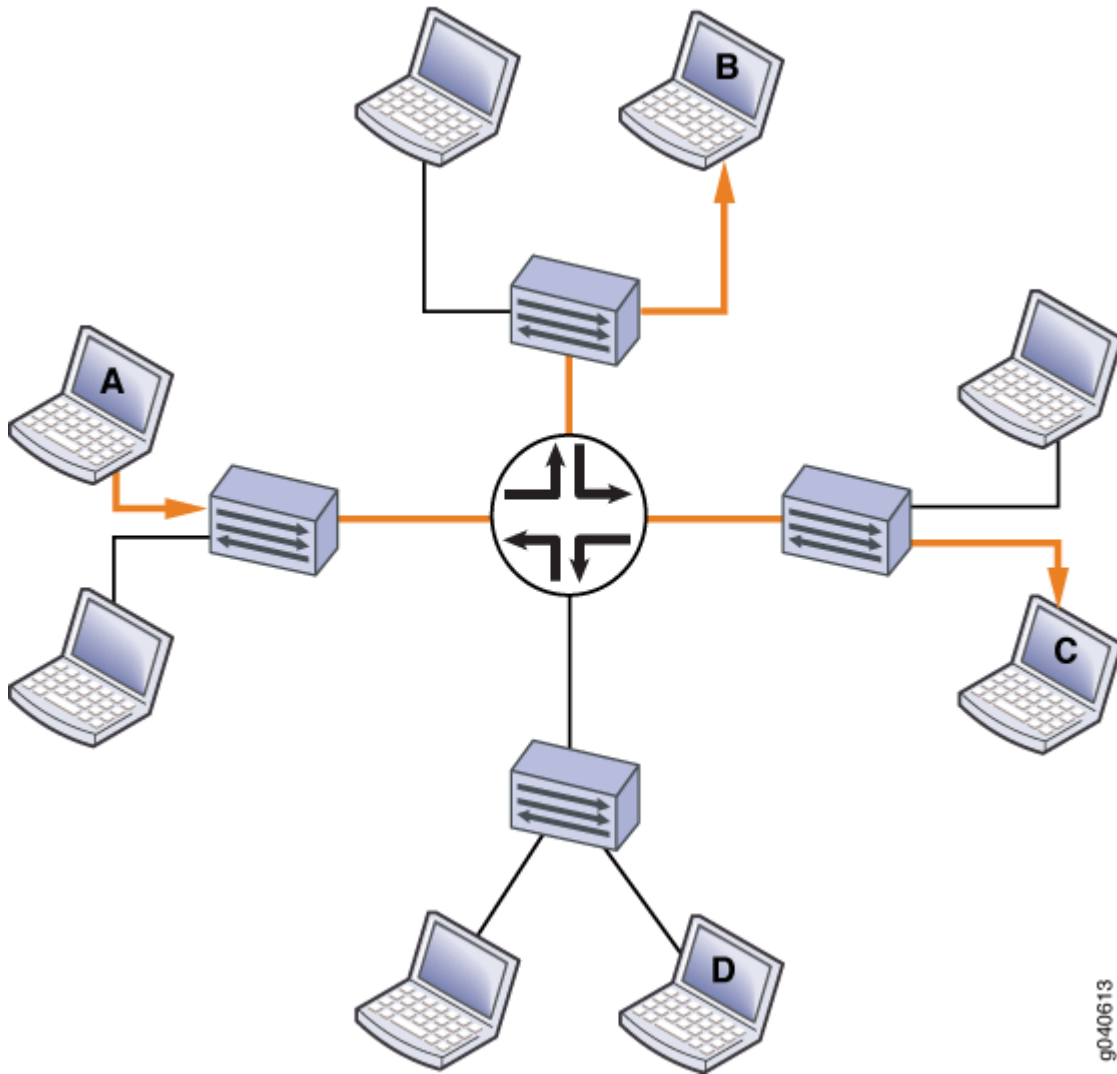


Figure 16 on page 154 shows the same networks with IGMP snooping configured. The Layer 2 devices forward multicast traffic to registered receivers only.

Figure 16: Networks with IGMP Snooping Configured



Configuration

IN THIS SECTION

- Procedure | 155

Procedure

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.

```
set bridge-domains domain1 domain-type bridge
set bridge-domains domain1 interface ge-0/0/1.1
set bridge-domains domain1 interface ge-0/0/2.1
set bridge-domains domain1 interface ge-0/0/3.1
set bridge-domains domain1 protocols igmp-snooping query-interval 200
set bridge-domains domain1 protocols igmp-snooping query-response-interval 0.4
set bridge-domains domain1 protocols igmp-snooping query-last-member-interval 0.1
set bridge-domains domain1 protocols igmp-snooping robust-count 4
set bridge-domains domain1 protocols igmp-snooping immediate-leave
set bridge-domains domain1 protocols igmp-snooping proxy
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/1.1 host-only-interface
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/1.1 group-limit 50
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/3.1 static group
225.100.100.100
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/2.1 multicast-router-
interface
```

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 [Junos OS CLI User Guide](#).

To configure IGMP snooping:

1. Configure the bridge domain.

```
[edit bridge-domains domain1]
user@host# set domain-type bridge
user@host# set interface ge-0/0/1.1
user@host# set interface ge-0/0/2.1
user@host# set interface ge-0/0/3.1
```

2. Enable IGMP snooping and configure the router to serve as a proxy.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping proxy
```

3. Configure the limit for the number of multicast groups allowed on the **ge-0/0/1.1** interface to 50.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/1.1 group-limit 50
```

4. Configure the router to immediately remove a group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping immediate-leave
```

5. Statically configure IGMP group membership on a port.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/3.1 static group 225.100.100.100
```

6. Configure an interface to be an exclusively router-facing interface (to receive multicast traffic).

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/2.1 multicast-router-interface
```

7. Configure an interface to be an exclusively host-facing interface (to drop IGMP query messages).

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/1.1 host-only-interface
```

8. Configure the IGMP message intervals and robustness count.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping robust-count 4
user@host# set protocols igmp-snooping query-last-member-interval 0.1
```

```
user@host# set protocols igmp-snooping query-interval 200
user@host# set protocols igmp-snooping query-response-interval 0.4
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show bridge-domains` command.

```
user@host# show bridge-domains
domain1 {
  domain-type bridge;
  interface ge-0/0/1.1;
  interface ge-0/0/2.1;
  interface ge-0/0/3.1;
  protocols {
    igmp-snooping {
      query-interval 200;
      query-response-interval 0.4;
      query-last-member-interval 0.1;
      robust-count 4;
      immediate-leave;
      proxy;
      interface ge-0/0/1.1 {
        host-only-interface;
        group-limit 50;
      }
      interface ge-0/0/3.1 {
        static {
          group 225.100.100.100;
        }
      }
      interface ge-0/0/2.1 {
        multicast-router-interface;
      }
    }
  }
}
```

Verification

To verify the configuration, run the following commands:

- `show igmp snooping interface`
- `show igmp snooping membership`
- `show igmp snooping statistics`

SEE ALSO

[Understanding IGMP Snooping | 142](#)

[Host-Side Interfaces and IGMP Snooping Proxy Mode | 146](#)

[Multicast-Router Interfaces and IGMP Snooping Proxy Mode | 145](#)

Configuring IGMP Snooping Trace Operations

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags.



NOTE: Enabling **traceoptions** under the `igmp-snooping` hierarchy enables **igmp-snooping**. Correspondingly, disabling **traceoptions** under the `igmp-snooping` hierarchy disables **igmp-snooping**.

The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
client-notification	Trace notifications.
general	Trace general flow.
group	Trace group operations.

(Continued)

Flag	Description
host-notification	Trace host notifications.
leave	Trace leave group messages (IGMPv2 only).
normal	Trace normal events.
packets	Trace all IGMP packets.
policy	Trace policy processing.
query	Trace IGMP membership query messages.
report	Trace membership report messages.
route	Trace routing information.
state	Trace state transitions.
task	Trace routing protocol task processing.
timer	Trace timer processing.

You can configure tracing operations for IGMP snooping globally or in a routing instance. The following example shows the global configuration.

To configure tracing operations for IGMP snooping:

1. Configure the filename for the trace file.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set file igmp-snoop-trace
```

2. (Optional) Configure the maximum number of trace files.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set file files 5
```

3. (Optional) Configure the maximum size of each trace file.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set file size 1m
```

4. (Optional) Enable unrestricted file access.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set file world-readable
```

5. Configure tracing flags. Suppose you are troubleshooting issues with a policy related to received packets on a particular logical interface with an IP address of 192.168.0.1. The following example shows how to flag all policy events for received packets associated with the IP address.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set flag policy receive | match 192.168.0.1
```

6. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-snoop-trace
```

SEE ALSO

[Tracing and Logging Junos OS Operations](#)

[Configuring IGMP Snooping | 147](#)

Configuring IGMP or MLD Snooping Version

IN THIS SECTION

- [Overview | 161](#)
- [Non-proxy mode | 162](#)
- [Proxy mode | 162](#)

Overview

Per RFC 4541, snooping by default is enabled in the IGMPv3 mode on interfaces that are a part of the VLAN or bridge-domain involved in snooping. MLD snooping defaults to MLDv2. When the router connected to the CPE sends IGMPv3 or MLDv2 general queries to refresh the IGMP group information, a standard CPE or end host that works as per the RFC should be able to process and respond to an IGMP v2/v3 or MLDv1/v2 query.

End hosts or CPE devices that are not compliant with RFC 4541 do not respond to IGMPv3 or MLDv2 queries sent from the L2 device enabled for snooping. This version mismatch could result in a traffic outage.

You can explicitly specify the IGMP or MLD snooping version for a VLAN or bridge-domain associated with L2 multicast. This makes sure that the router sends an IGMP or MLD startup query or periodic general query (when configured for L2 querier) with the configured version of IGMP or MLD snooping. This in turn ensures that the non-complaint CPE devices respond with the IGMP or MLD group/report information. This helps in keeping the multicast stream alive with hosts and CPEs that can't respond to IGMP or MLD queries of a higher version.



NOTE: IGMP or MLD snooping version is configured on a VLAN or bridge-domain and is applicable to all interfaces (IFLs) under that VLAN or bridge-domain.

The IGMP/MLD snooping version configuration for VLAN or bridge-domain can be done in two ways, the default routing instance (global instance) or in a specific routing instance.

With version configuration, MCSNOOPD (the multicast snooping process that allows Layer 3 inspection from Layer 2 device) refreshes reports from the hosts using:

- **Startup Query:** This is the general IGMP query sent whenever snooping is enabled in a VLAN or bridge-domain. This query is sent when MCSNOOPD builds the interface state of a member of VLAN or bridge-domain for the first time or when the interface (IFL or IFD) of a VLAN or bridge-domain flaps.

- L2 Querier: This is a periodic feature that can be explicitly enabled in a specific VLAN or bridge-domain enabled for snooping.
- Proxy: IGMP or MLD proxy reports on receiving IGMP or MLD query packets from the remote querier.



NOTE: In EVPN deployments this snooping version configuration is applicable to access facing interfaces only.

Non-proxy mode

Without a proxy arrangement, IGMP snooping does not generate or introduce queries and reports. It will only snoop reports received from all of its interfaces (including multicast-router interfaces) to build its state and group (S,G) database.

- Startup queries are generated when the interface flaps or when a new member of VLAN or bridge-domain is added. The startup queries are general IGMP or MLD query messages (source IP Address=0.0.0.0) and will be sent over the interface with the configured version or it defaults to IGMPv3 or MLD v2.
- If the L2 querier feature is enabled for a VLAN or bridge-domain, then periodic queries (default timer is 125secs) are generated and sent to all the VLAN or bridge-domain members with the configured version or it defaults to the IGMPv3 or MLDv2 query.

Proxy mode

In this mode, Proxy snooping reduces the number of IGMP reports sent towards an IGMP router.

- Startup queries behave in the same way as in non-proxy mode.
- With the L2 querier enabled, it behaves in the same way as in non-proxy mode.
- The version used for IGMP reports generated in response to a general query received from a peer is based on the host compatibility mode and is decided as below:
 - Without version configuration, on receiving a query with a lower version, the mode is set to the lowest version and the corresponding querier-present timer is started. On expiry of the querier-present timer the version changes to IGMPv3 or MLDv2, based on the protocol the interface is operating on.
 - With version configuration, on receiving a query with a lower version, the mode is set to the lowest version and the corresponding querier-present timer is started. On expiry of the querier-present timer the version updates to the configured version as per the version set by the config

statement set routing-instance <routing-instance-name> protocols igmp-snooping version <version> for example.

- The group specific query generated in response to a leave is based on the lowest version of the IGMP or MLD report and the configured interface version. If the configured interface version is lower, then that version is used for group-compatibility.

SEE ALSO

version (IGMP Snooping)

version (MLD Snooping)

Restrict Flooding of Unknown Multicast Traffic to the Multicast-router Interface in an L2 environment

SUMMARY

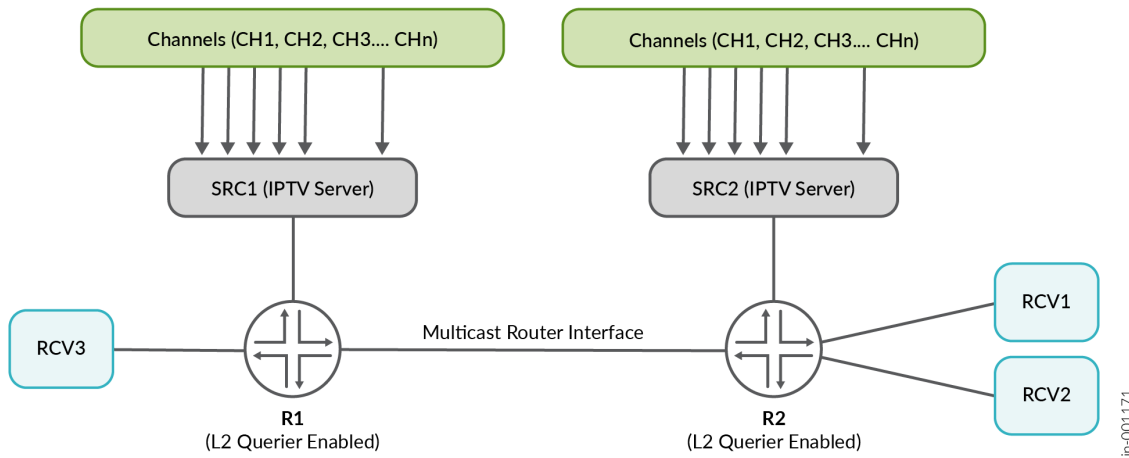
Limit the flooding of L2 multicast streams to multicast-router interfaces on a per VLAN or bridge domain level. By restricting the flooding of multicast data to the multicast-router interface based on IGMP or MLD join messages, unknown multicast traffic can be limited, resulting in efficient bandwidth utilization.

IN THIS SECTION

- [Restricting flooding of unknown multicast traffic in an EVPN environment | 165](#)

Let's take an example of an L2 multicast video scenario where routers R1 and R2 are enabled with IGMP/MLD Snooping.

Figure 17: L2 Multicast Video (IPTV)



SRC1 and SRC2 are IPTV servers with a video headend, generating multiple IPTV channels - CH1, CH2 and CH3. The L2 querier is enabled on routers R1 and R2 to ensure that the hosts can refresh reports periodically, which will time out and remove the receivers from multicast forwarding otherwise.

In the topology above, router R2 has receivers interested in receiving channels CH1 and CH2 but will receive multicast streams from CH3 too. This is by virtue of it being a multicast-router interface.

Similarly, router R1 has a receiver interested only in channel CH3 but will receive unsolicited multicast streams from CH1 and CH2 over the multicast-router interface.

By configuring `no-flood-to-multicast-router-interfaces`, you can limit the above-mentioned bandwidth inefficiency effectively by sending multicast streams only when there are interested receivers behind a multicast-router interface.



NOTE: This feature is specific to an L2 multicast environment.

If PIM-enabled L3 routers (with IGMP/MLD querier turned on implicitly) are present in the topology and `no-flood-to-multicast-router-interfaces` is enabled, then it may result in traffic outage for receivers behind the L3 multicast routers. This is because multicast routers typically send PIM join messages on seeing an IGMP join from a last hop router (LHR).

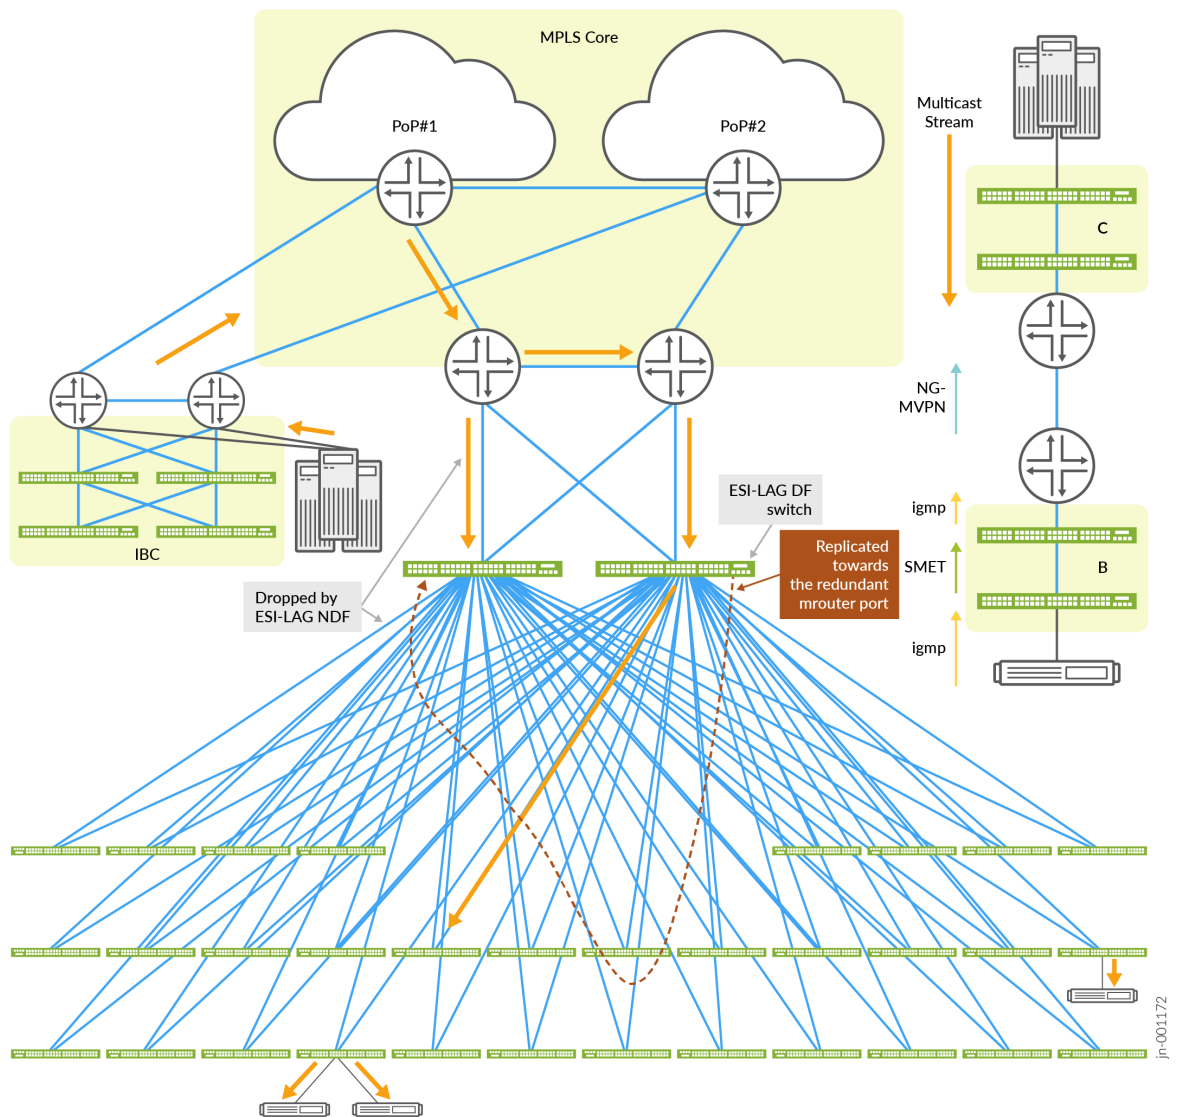
Based on your network design, `no-flood-to-multicast-router-interfaces` must be used judiciously.

Restricting flooding of unknown multicast traffic in an EVPN environment

Let us take an example of a typical EVPN topology which is an EVPNoVXLAN bridge overlay with SMET (Selective multicast forwarding) enabled. EVPN type 6 routes are advertised, ensuring forwarding of multicast traffic to interested receivers only.

There are IBC headend devices which are connected to a redundant external multicast-router and the EVPN fabric below. Since this is a bridge overlay there is only L2 connectivity from the redundant external multicast routers where IRB is enabled without snooping. The multicast traffic coming into the external multicast routers are bridged or L2 switched when traffic needs to egress from the IRB.

Figure 18: EVPN Multicast deployment with Snooping



With SMET bridged overlay fabric stitched to the redundant external multicast-router ports, multicast traffic entering one multicast-router port is circled via the fabric access switches back to the second multicast-router port (the orange dotted line in [Figure 18 on page 166](#)). The EVPN PE device attracts all multicast traffic, known and unknown over the EVPN core. Typically, EVPN PEs send EVPN Type 3 routes that have multicast extended community flags without the IGMP/MLD snooping proxy set in its NLRI. By setting the `no-flood-to-multicast-router-interfaces` configuration statement, the EVPN Type 3 routes are sent with multicast extended community flags with the IGMP/MLD snooping proxy flags set. As a result, the multicast-router interface will not attract unknown multicast traffic.



NOTE: This feature is typically implemented in an EVPN bridge overlay scenario. If L3 routing is operational across the external multicast router, enabling `no-flood-to-multicast-router-interfaces` may result in traffic outage.

SEE ALSO

No Link Title

RELATED DOCUMENTATION

[Example: Configuring IGMP Snooping | 141](#)

Example: Configuring IGMP Snooping on SRX Series Devices

IN THIS SECTION

- [Requirements | 167](#)
- [Overview and Topology | 168](#)
- [Configuration | 169](#)
- [Verifying IGMP Snooping Operation | 173](#)

You can enable IGMP snooping on a VLAN to constrain the flooding of IPv4 multicast traffic on a VLAN. When IGMP snooping is enabled, the device examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the device then forwards multicast traffic only to those interfaces that are connected to relevant receivers instead of flooding the traffic to all interfaces.

This example describes how to configure IGMP snooping:

Requirements

This example uses the following hardware and software components:

- One SRX Series Firewall

- Junos OS Release 18.1R1

Before you configure IGMP snooping, be sure you have:

- Configured a VLAN, v1, on the device
- Assigned interfaces ge-0/0/1, ge-0/0/2, ge-0/0/3, and ge-0/0/4 to v1
- Configured ge-0/0/3 as a trunk interface

Overview and Topology

IN THIS SECTION

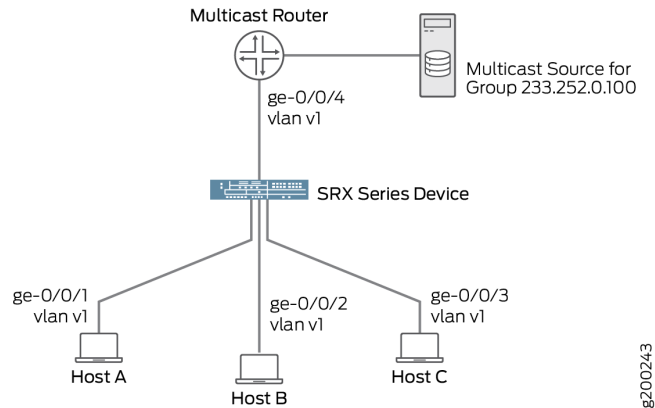
- [Topology | 168](#)

IGMP snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the SRX Series Firewall broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, the SRX Series Firewall monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the SRX Series Firewall to keep track of the multicast groups and associated member ports. The SRX Series Firewall uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts.

Topology

The sample topology is illustrated in [Figure 19 on page 169](#).

Figure 19: IGMP Snooping Sample Topology



In this sample topology, the multicast router forwards multicast traffic to the device from the source when it receives a membership report for group 233.252.0.100 from one of the hosts—for example, Host B. If IGMP snooping is not enabled on vlan100, the device floods the multicast traffic on all interfaces in vlan100 (except for interface ge-0/0/2.0). If IGMP snooping is enabled on vlan100, the device monitors the IGMP messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The device then forwards the multicast traffic only to interface ge-0/0/2.

Configuration

IN THIS SECTION

- [Procedure | 170](#)

To configure IGMP snooping on a device:

Procedure

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.

```
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members v1
set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members v1
set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members v1
set interfaces ge-0/0/4 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/4 unit 0 family ethernet-switching vlan members v1
set vlans v1 vlan-id 100
set protocols igmp-snooping vlan v1 query-interval 200
set protocols igmp-snooping vlan v1 query-response-interval 0.4
set protocols igmp-snooping vlan v1 query-last-member-interval 0.1
set protocols igmp-snooping vlan v1 robust-count 4
set protocols igmp-snooping vlan v1 immediate-leave
set protocols igmp-snooping vlan v1 proxy
set protocols igmp-snooping vlan v1 interface ge-0/0/1.0 host-only-interface
set protocols igmp-snooping vlan v1 interface ge-0/0/1.0 group-limit 50
set protocols igmp-snooping vlan v1 interface ge-0/0/4.0 static group 233.252.0.100
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IGMP snooping:

1. Configure the access mode interfaces.

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode access
user@host# set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members v1
user@host# set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode access
user@host# set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members v1
```

```

user@host# set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode trunk
user@host# set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members v1
user@host# set interfaces ge-0/0/4 unit 0 family ethernet-switching interface-mode access
user@host# set interfaces ge-0/0/4 unit 0 family ethernet-switching vlan members v1

```

2. Configure the VLAN.

```

[edit]
user@host# set vlans v1 vlan-id 100

```

3. Enable IGMP snooping and configure the device to serve as a proxy.

```

[edit]
user@host# set protocols igmp-snooping vlan v1 proxy

```

4. Configure the limit for the number of multicast groups allowed on the ge-0/0/1.0 interface to 50.

```

[edit]
user@host# set protocols igmp-snooping vlan v1 interface ge-0/0/1.0 group-limit 50

```

5. Configure the device to immediately remove a group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged.

```

[edit]
user@host# set protocols igmp-snooping vlan v1 immediate-leave

```

6. Statically configure interface ge-0/0/4 as a multicast-router interface.

```

[edit]
user@host# set protocols igmp-snooping vlan v1 interface ge-0/0/4.0 static group 233.252.0.100

```

7. Configure an interface to be an exclusively host-facing interface (to drop IGMP query messages).

```

[edit]
user@host# set protocols igmp-snooping vlan v1 interface ge-0/0/1.0 host-only-interface

```

8. Configure the IGMP message intervals and robustness count.

```
[edit]
user@host# set protocols igmp-snooping vlan v1 query-interval 200
user@host# set protocols igmp-snooping vlan v1 query-response-interval 0.4
user@host# set protocols igmp-snooping vlan v1 query-last-member-interval 0.1
user@host# set protocols igmp-snooping vlan v1 robust-count 4
```

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

```
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the `show protocols igmp-snooping` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show protocols igmp-snooping
  vlan v1 {
    query-interval 200;
    query-response-interval 0.4;
    query-last-member-interval 0.1;
    robust-count 4;
    immediate-leave;
    proxy;
    interface ge-0/0/1.0 {
      host-only-interface;
      group-limit 50;
    }
    interface ge-0/0/4.0 {
      static {
        group 233.252.0.100;
      }
    }
  }
}
```

Verifying IGMP Snooping Operation

IN THIS SECTION

- [Displaying IGMP Snooping Information for VLAN v1 | 173](#)

To verify that IGMP snooping is operating as configured, perform the following task:

Displaying IGMP Snooping Information for VLAN v1

Purpose

Verify that IGMP snooping is enabled on vlan v1 and that ge-0/0/4 is recognized as a multicast-router interface.

Action

From operational mode, enter the `show igmp snooping membership` command.

```
user@host> show igmp snooping membership
Instance: default-switch

Vlan: v1

Learning-Domain: default
Interface: ge-0/0/4.0, Groups: 1
Group: 233.252.0.100
Group mode: Exclude
Source: 0.0.0.0
Last reported by: Local
Group timeout: 0 Type: Static
```

Meaning

By showing information for vlanv1, the command output confirms that IGMP snooping is configured on the VLAN. Interface ge-0/0/4.0 is listed as a multicast-router interface, as configured. Because none of the host interfaces are listed, none of the hosts are currently receivers for the multicast group.

RELATED DOCUMENTATION

[IGMP Snooping Overview | 96](#)

igmp-snooping

Configuring Point-to-Multipoint LSP with IGMP Snooping

By default, IGMP snooping in VPLS uses multiple parallel streams when forwarding multicast traffic to PE routers participating in the VPLS. However, you can enable point-to-multipoint LSP for IGMP snooping to have multicast data traffic in the core take the point-to-multipoint path rather than using a pseudowire path. The effect is a reduction in the amount of traffic generated on the PE router when sending multicast packets for multiple VPLS sessions.

Figure 1 shows the effect on multicast traffic generated on the PE1 router (the device where the setting is enabled). When pseudowire LSP is used, the PE1 router sends multiple packets whereas with point-to-multipoint LSP enabled, only a single copy of the packets on the PE1 router is sent.

The options configured for IGMP snooping are applied on a per routing-instance, so all IGMP snooping routes in the same instance will use the same mode, point-to-multipoint or pseudowire.

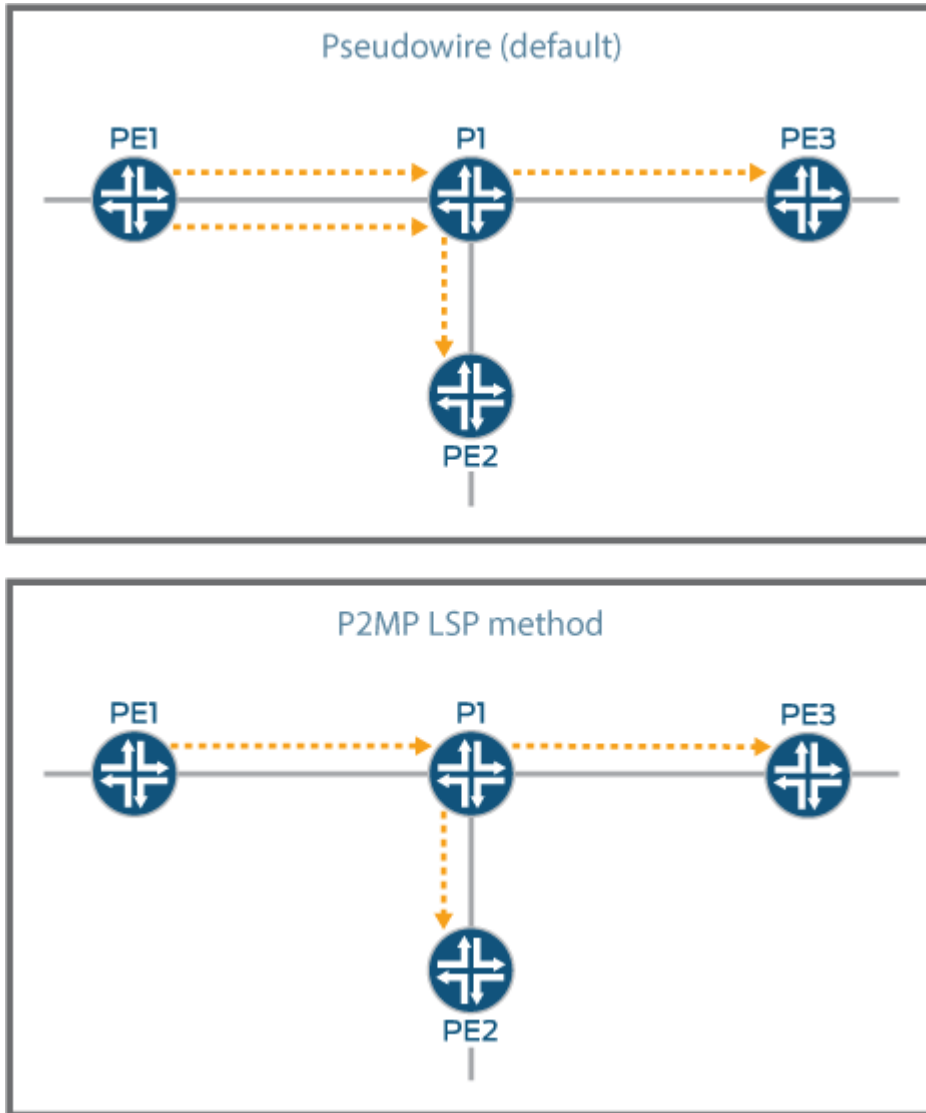


NOTE: The point-to-multipoint option is available on MX960, MX480, MX240, and MX80 routers running Junos OS 13.3 and later.



NOTE: IGMP snooping is not supported on the core-facing pseudowire interfaces; all PE routers participating in VPLS will continue to receive multicast data traffic even when this option is enabled.

Figure 20: Point-to-multipoint LSP generates less traffic on the PE router than pseudowire.



In a VPLS instance with IGMP-snooping that uses a point-to-multipoint LSP, mcsnooped (the multicast snooping process that allows Layer 3 inspection from Layer 2 device) will start listening for point-to-multipoint next-hop notifications and then manage the IGMP snooping routes accordingly. Enabling the `use-p2mp-lsp` command in Junos allows the IGMP snooping routes to start using this next-hop. In short, if point-to-multipoint is configured for a VPLS instance, multicast data traffic in the core can avoid ingress replication by taking the point-to-multipoint path. If the point-to-multipoint next-hop is unavailable, packets are handled in the VPLS instance in the same way as broadcast packets or unknown unicast frames. Note that IGMP snooping is not supported on the core-facing pseudowire interfaces. PE routers participating in VPLS will continue to receive multicast data traffic regardless of how Point-to-Multipoint is set.

To enable point-to-multipoint LSP, type the following CLI command:

```
[edit]
user@host> set routing-instances instance name instance-type vpls igmp-snooping-options use-p2mp-
lsp
```

The following output shows the hierarchical presence of `igmp-snooping-options`:

```
routing-instances {
  <instance-name> {
    instance-type vpls;
    igmp-snooping-options {
      use-p2mp-lsp;
    }
  }
}
```

To show the operational status of point-to-multipoint LSP for IGMP snooping routes, use the following CLI command:

```
user@host> show igmp snooping options
```

```
Instance: master
  P2MP LSP in use: no
Instance: default-switch
  P2MP LSP in use: no
Instance: name
  P2MP LSP in use: yes
```

RELATED DOCUMENTATION

use-p2mp-lsp

show igmp snooping options

multicast-snooping-options

Configuring MLD Snooping

IN THIS CHAPTER

- [Understanding MLD Snooping | 177](#)
- [Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)
- [Configuring MLD Snooping on a Switch VLAN with ELS Support \(CLI Procedure\) | 197](#)
- [Example: Configuring MLD Snooping on EX Series Switches | 204](#)
- [Example: Configuring MLD Snooping on SRX Series Devices | 209](#)
- [Configuring MLD Snooping Tracing Operations on EX Series Switches \(CLI Procedure\) | 217](#)
- [Configuring MLD Snooping Tracing Operations on EX Series Switch VLANs \(CLI Procedure\) | 220](#)
- [Example: Configuring MLD Snooping on EX Series Switches | 223](#)
- [Example: Configuring MLD Snooping on Switches with ELS Support | 228](#)
- [Verifying MLD Snooping on EX Series Switches \(CLI Procedure\) | 234](#)
- [Verifying MLD Snooping on Switches | 239](#)

Understanding MLD Snooping

IN THIS SECTION

- [Benefits of MLD Snooping | 178](#)
- [How MLD Snooping Works | 178](#)
- [MLD Message Types | 179](#)
- [How Hosts Join and Leave Multicast Groups | 180](#)
- [Support for MLDv2 Multicast Sources | 181](#)
- [MLD Snooping and Forwarding Interfaces | 181](#)
- [General Forwarding Rules | 182](#)
- [Examples of MLD Snooping Multicast Forwarding | 183](#)

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs. When MLD snooping is enabled on a VLAN, a Juniper Networks device examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. On the basis of what it learns, the device then forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces.

MLD snooping supports MLD version 1 (MLDv1) and MLDv2. For details on MLDv1 and MLDv2, see the following standards:

- MLDv1—See RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*.
- MLDv2—See RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*.

Benefits of MLD Snooping

- **Optimized bandwidth utilization**—The main benefit of MLD snooping is to reduce flooding of packets. IPv6 multicast data is selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN.
- **Improved security**—Denial of service attacks from unknown sources are prevented.


How MLD Snooping Works

By default, the device floods Layer 2 multicast traffic on all of the interfaces belonging to that VLAN on the device, except for the interface that is the source of the multicast traffic. This behavior can consume significant amounts of bandwidth.

You can enable MLD snooping to avoid this flooding. When you enable MLD snooping, the device monitors MLD messages between receivers (hosts) and multicast routers and uses the content of the messages to build an IPv6 multicast forwarding table—a database of IPv6 multicast groups and the interfaces that are connected to the interested members of each group. When the device receives multicast traffic for a multicast group, it uses the forwarding table to forward the traffic only to interfaces that are connected to receivers that belong to the multicast group.

[Figure 21 on page 179](#) shows an example of multicast traffic flow with MLD snooping enabled.

Figure 21: Multicast Traffic Flow with MLD Snooping Enabled



MLD Message Types

Multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. In any given subnet, one multicast router is elected to act as an MLD querier. The MLD querier sends out the following types of queries to hosts:

- General query—Asks whether any host is listening to any group.
- Group-specific query—Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.
- Group-and-source-specific query—(MLD version 2 only) Asks whether any host is listening to group multicast traffic from a specific multicast source. This query is sent in response to a host indicating that it is no longer interested in receiving group multicast traffic from the multicast source and allows the router to quickly determine any remaining hosts are interested in receiving group multicast traffic from that source.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report—Indicates that the host wants to join a particular multicast group.
- Leave report—Indicates that the host wants to leave a particular multicast group.

Only MLDv1 hosts use two different kinds of reports to indicate whether they want to join or leave a group. MLDv2 hosts send only one kind of report, the contents of which indicate whether they want to join or leave a group. However, for simplicity's sake, the MLD snooping documentation uses the term *membership report* for a report that indicates that a host wants to join a group and uses the term *leave report* for a report that indicates a host wants to leave a group.

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

- By sending an unsolicited membership report that specifies the multicast group that the host is attempting to join.
- By sending a membership report in response to a query from a multicast router.

A multicast router continues to forward multicast traffic to an interface provided that at least one host on that interface responds to the periodic general queries indicating its membership. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general queries.

Hosts can leave multicast groups in either of two ways:

- By not responding to periodic queries within a set interval of time. This results in what is known as a "silent leave."
- By sending a leave report.



NOTE: If a host is connected to the device through a hub, the host does not automatically leave the multicast group if it disconnects from the hub. The host remains a member of the group until group membership times out and a silent leave occurs. If another host connects to the hub port before the silent leave occurs, the new host might receive the group multicast traffic until the silent leave, even though it never sent an membership report.

Support for MLDv2 Multicast Sources

In MLDv2, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source *except* the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources.

Devices that support MLD snooping support MLDv2 membership reports that are in INCLUDE and EXCLUDE mode. However, SRX Series Firewalls, QFX Series switches, and EX Series switches running MLD snooping, except for EX9200 switches, do not support forwarding on a per-source basis. Instead, the device consolidates all INCLUDE and EXCLUDE mode reports it receives on a VLAN for a specified group into a single route that includes all multicast sources for that group, with the next hop being all interfaces that have interested receivers for the group. As a result, interested receivers on the VLAN can receive traffic from a source that they did not include in their INCLUDE report or from a source they excluded in their EXCLUDE report. For example, if Host 1 wants traffic for group G from Source A and Host 2 wants traffic for group G from Source B, they both receive traffic for group G regardless of whether A or B sends the traffic.

MLD Snooping and Forwarding Interfaces

To determine how to forward multicast traffic, the device with MLD snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or MLD queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The device learns about these interfaces by monitoring MLD traffic. If an interface receives MLD queries, the device adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the device adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the device learns about are subject to aging. For example, if a learned multicast-router interface does not receive MLD queries within a certain interval, the device removes the entry for that interface from its multicast forwarding table.



NOTE: For the device to learn multicast-router interfaces and group-member interfaces, an MLD querier must exist in the network. For the device itself to function as an MLD querier, MLD must be enabled on the device.

You can statically configure an interface to be a multicast-router interface or a group-member interface. The device adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. You can have a mix of statically configured and dynamically learned interfaces on the device.

General Forwarding Rules

Multicast traffic received on the device interface in a VLAN on which MLD snooping is enabled is forwarded according to the following rules.

MLD protocol traffic is forwarded as follows:

- MLD general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- MLD group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- MLD reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not MLD protocol traffic is forwarded as follows:

- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.



NOTE: When IGMP and MLD snooping are both enabled on the same VLAN, multicast-router interfaces are created as part of IGMP and MLD snooping configuration. Unregistered multicast traffic is not blocked and can be passed through router interfaces, so due to hardware limitations, unregistered IPv4 multicast traffic might be passed through the multicast router interfaces created as part of MLD snooping configuration,

and unregistered IPv6 multicast traffic might pass through multicast-router interfaces created as part of IGMP snooping configuration.

Examples of MLD Snooping Multicast Forwarding

The following examples are provided to illustrate how MLD snooping forwards multicast traffic in different topologies:

Scenario 1: Device Forwarding Multicast Traffic to a Multicast Router and Hosts

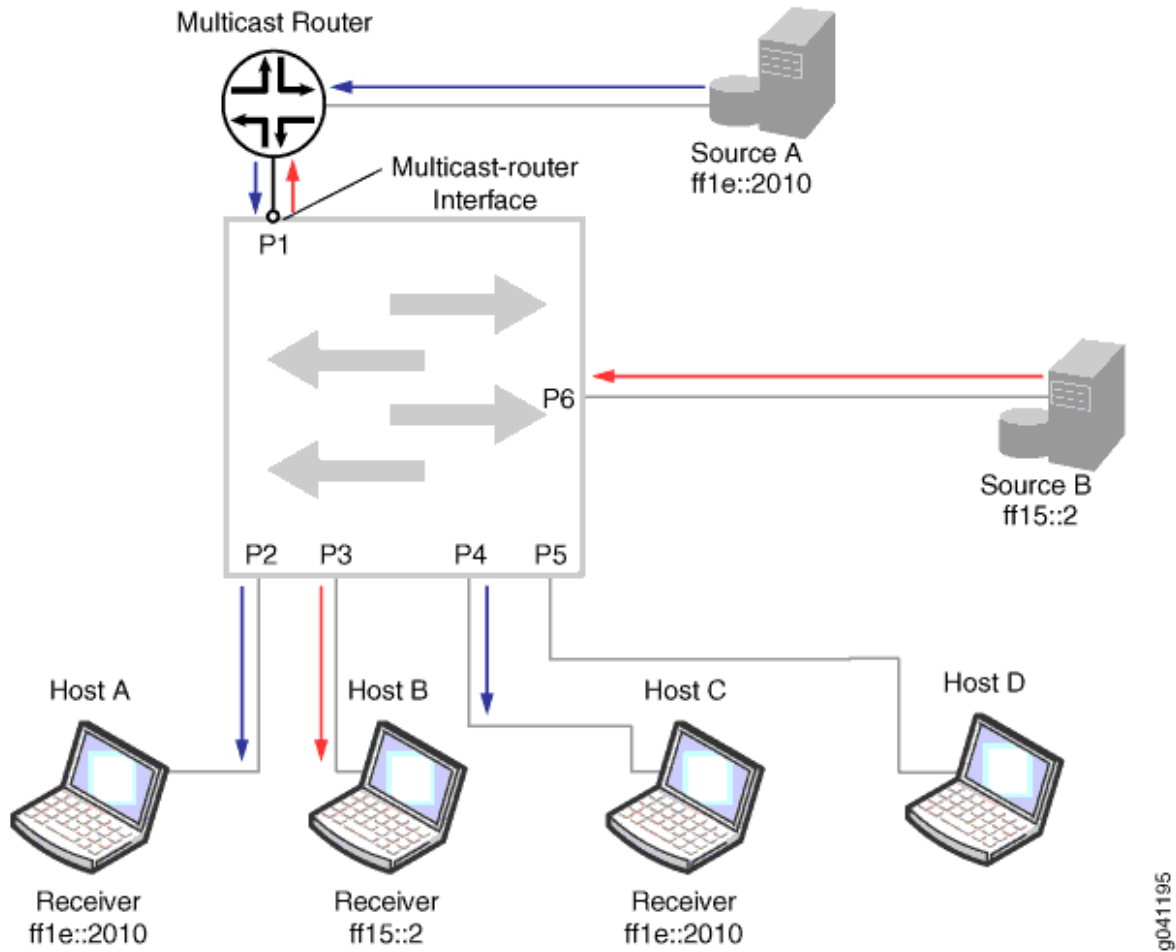
In the topology shown in [Figure 22 on page 184](#), the device acting as a Layer 2 device receives multicast traffic belonging to multicast group **ff1e::2010** from Source A, which is connected to the multicast router. It also receives multicast traffic belonging to multicast group **ff15::2** from Source B, which is connected directly to the device. All interfaces on the device belong to the same VLAN.

Because the device receives MLD queries from the multicast router on interface P1, MLD snooping learns that interface P1 is a multicast-router interface and adds the interface to its multicast forwarding table. It forwards any MLD general queries it receives on this interface to all host interfaces on the device, and, in turn, forwards membership reports it receives from hosts to the multicast-router interface.

In the example, Hosts A and C have responded to the general queries with membership reports for group **ff1e::2010**. MLD snooping adds interfaces P2 and P4 to its multicast forwarding table as member interfaces for group **ff1e::2010**. It forwards the group multicast traffic received from Source A to Hosts A and C, but not to Hosts B and D.

Host B has responded to the general queries with a membership report for group **ff15::2**. The device adds interface P3 to its multicast forwarding table as a member interface for group **ff15::2** and forwards multicast traffic it receives from Source B to Host B. The device also forwards the multicast traffic it receives from Source B to the multicast-router interface P1.

Figure 22: Scenario 1: Device Forwarding Multicast Traffic to a Multicast Router and Hosts

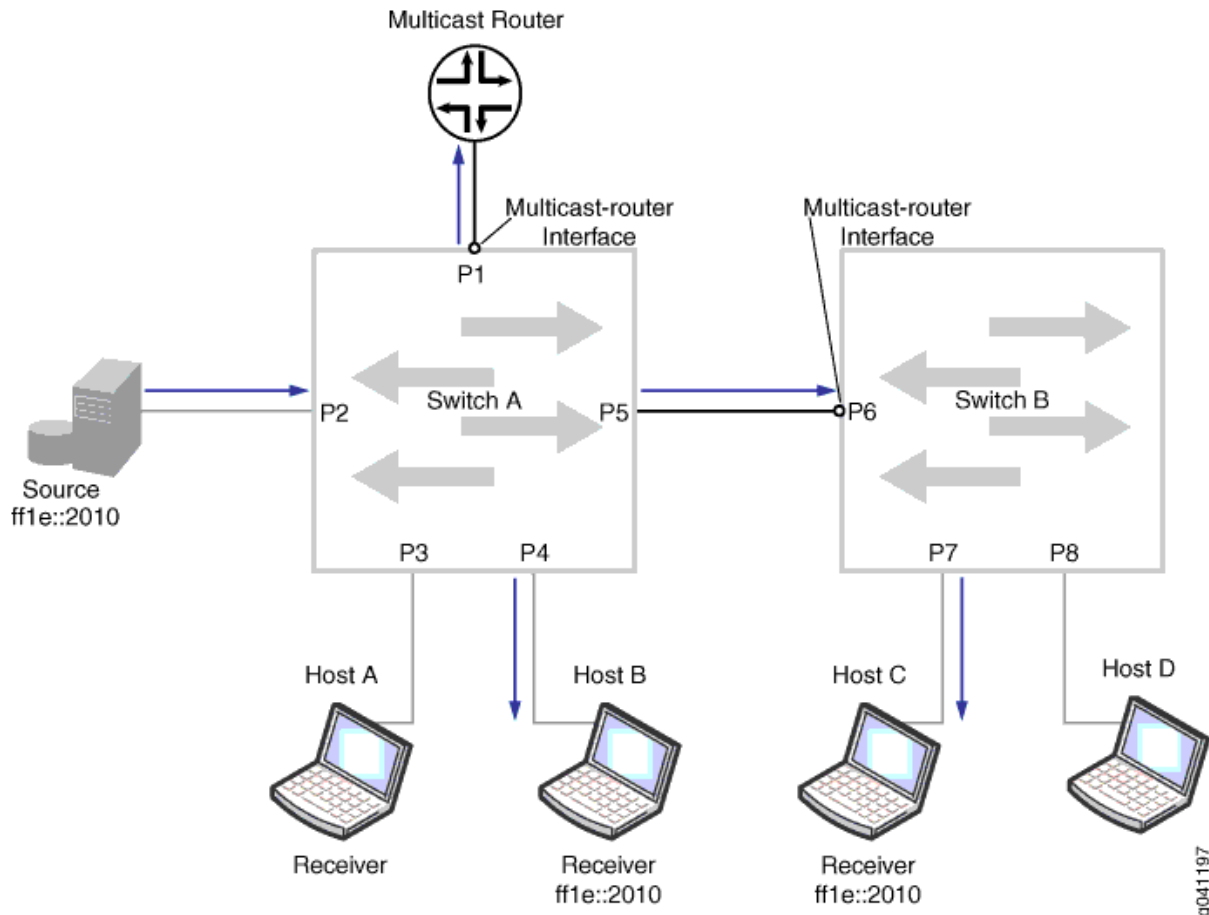


Scenario 2: Device Forwarding Multicast Traffic to Another Device

In the topology shown in [Figure 23 on page 185](#), a multicast source is connected to Device A. Device A in turn is connected to another device, Device B. Hosts on both Device A and B are potential members of the multicast group. Both devices are acting as Layer 2 devices, and all interfaces on the devices are members of the same VLAN.

Device A receives MLD queries from the multicast router on interface P1, making interface P1 a multicast-router interface for Device A. Device A forwards all general queries it receives on this interface to the other interfaces on the device, including the interface connecting Device B. Because Device B receives the forwarded MLD queries on interface P6, P6 is the multicast-router interface for Device B. Device B forwards the membership report it receives from Host C to Device A through its multicast-router interface. Device A forwards the membership report to its multicast-router interface, includes interface P5 in its multicast forwarding table as a group-member interface, and forwards multicast traffic from the source to Device B.

Figure 23: Scenario 2: Device Forwarding Multicast Traffic to Another Device



In certain implementations, you might have to configure P6 on Device B as a static multicast-router interface to avoid a delay in a host receiving multicast traffic. For example, if Device B receives unsolicited membership reports from its hosts before it learns which interface is its multicast-router interface, it does not forward those reports to Device A. If Device A then receives multicast traffic, it does not forward the traffic to Device B, because it has not received any membership reports on interface P5. This issue will resolve when the multicast router sends out its next general query; however, it can cause a delay in the host receiving multicast traffic. You can statically configure interface P6 as a multicast-router interface to solve this issue.

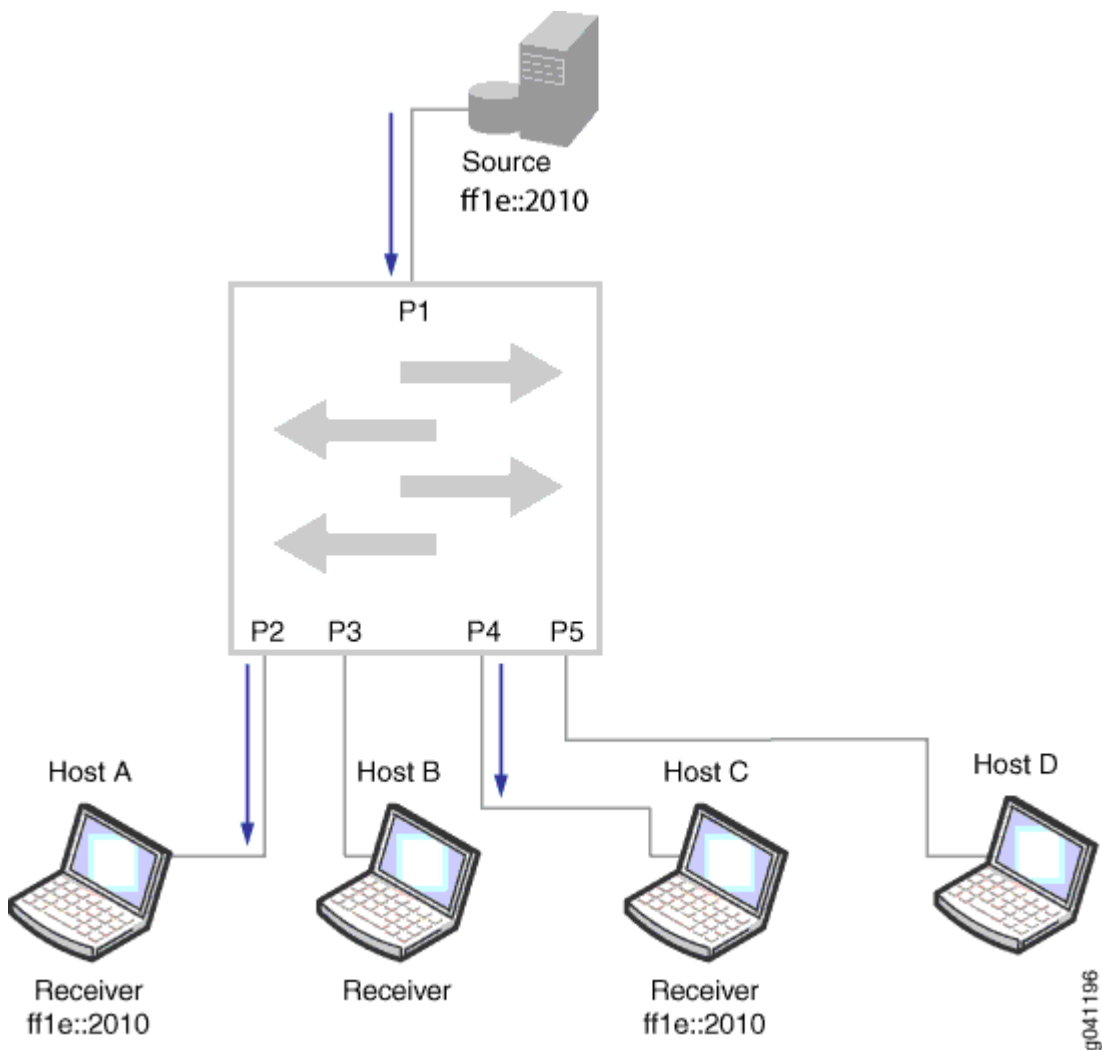
Scenario 3: Device Connected to Hosts Only (No MLD Querier)

In the topology shown in [Figure 24 on page 186](#), the device is connected to a multicast source and to hosts. There is no multicast router in this topology—hence there is no MLD querier. Without an MLD querier to respond to, a host does not send periodic membership reports. As a result, even if the host sends an unsolicited membership report to join a multicast group, its membership in the multicast group will time out.

For MLD snooping to work correctly in this network so that the device forwards multicast traffic to Hosts A and C only, you can either:

- Configure interfaces P2 and P4 as static group-member interfaces.
- Configure a *routed VLAN interface (RVI)*, also referred to as an integrated routing and bridging (IRB) interface, on the VLAN and enable MLD on it. In this case, the device itself acts as an MLD querier, and the hosts can dynamically join the multicast group and refresh their group membership by responding to the queries.

Figure 24: Scenario 3: Device Connected to Hosts Only (No MLD Querier)

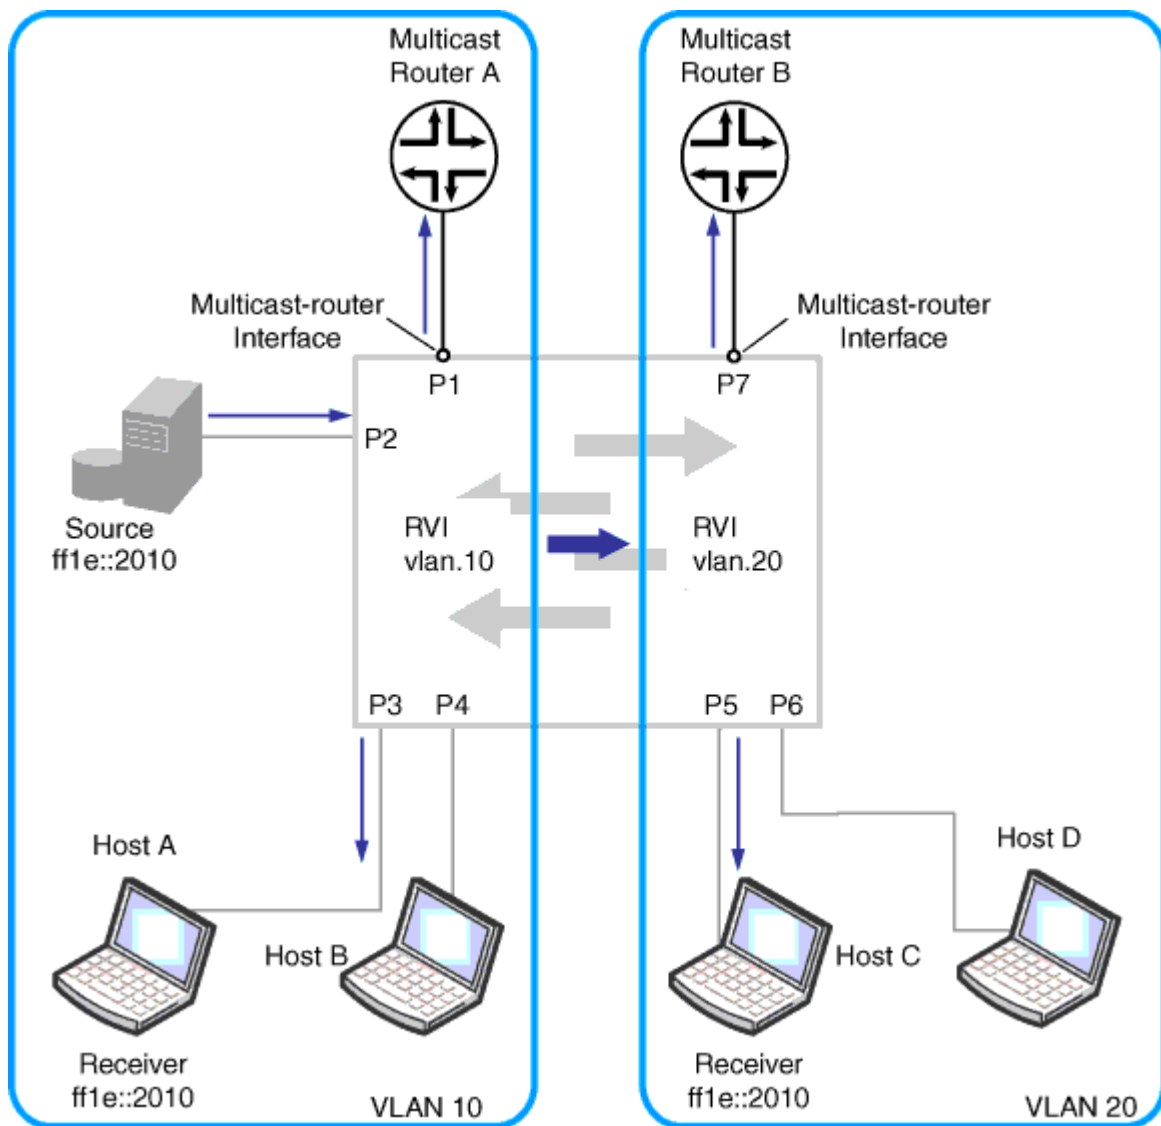


Scenario 4: Layer 2/Layer 3 Device Forwarding Multicast Traffic Between VLANs

In the topology shown in [Figure 25 on page 187](#), a multicast source, Multicast Router A, and Hosts A and B are connected to the device and are in VLAN 10. Multicast Router B and Hosts C and D are also connected to the device and are in VLAN 20.

In a pure Layer 2 environment, traffic is not forwarded between VLANs. For Host C to receive the multicast traffic from the source on VLAN 10, RVIs (or IRB interfaces) must be created on VLAN 10 and VLAN 20 to permit routing of the multicast traffic between the VLANs.

Figure 25: Scenario 4: Layer 2/Layer 3 device Forwarding Multicast Traffic Between VLANs



RELATED DOCUMENTATION

[Example: Configuring MLD Snooping on SRX Series Devices | 209](#)

[Configuring MLD Snooping on a Switch VLAN with ELS Support \(CLI Procedure\) | 197](#)

[Example: Configuring MLD Snooping on Switches with ELS Support | 228](#)

[Verifying MLD Snooping on Switches | 239](#)

[Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)

[Example: Configuring MLD Snooping on EX Series Switches | 204](#)

[Verifying MLD Snooping on EX Series Switches \(CLI Procedure\) | 234](#)

Configuring MLD Snooping on an EX Series Switch VLAN (CLI Procedure)

IN THIS SECTION

- [Enabling or Disabling MLD Snooping on VLANs | 190](#)
- [Configuring the MLD Version | 191](#)
- [Enabling Immediate Leave | 192](#)
- [Configuring an Interface as a Multicast-Router Interface | 193](#)
- [Configuring Static Group Membership on an Interface | 194](#)
- [Changing the Timer and Counter Values | 195](#)

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

MLD snooping is not enabled on the switch by default. To enable MLD snooping on all VLANs:

```
[edit]  
user@switch# set protocols mld-snooping vlan all
```

For many networks, MLD snooping requires no further configuration.

You can perform the following optional configurations per VLAN:

- Selectively enable MLD snooping on specific VLANs.



NOTE: You cannot configure MLD snooping on a secondary VLAN.

- Specify the MLD version for the general query that the switch sends on an interface when the interface comes up.
- Enable immediate leave on a VLAN or all VLANs. Immediate leave reduces the length of time it takes the switch to stop forwarding multicast traffic when the last member host on the interface leaves the group.
- Configure an interface as a static multicast-router interface for a VLAN or for all VLANs so that the switch does not need to dynamically learn that the interface is a multicast-router interface.
- Configure an interface as a static member of a multicast group so that the switch does not need to dynamically learn the interface's membership.
- Change the value for certain timers and counters to match the values configured on the multicast router serving as the MLD querier.



TIP: When you configure MLD snooping using the `vlan all` statement, any VLAN that is not individually configured for MLD snooping inherits the **vlan all** configuration. Any VLAN that is individually configured for MLD snooping, on the other hand, inherits none of its configuration from **vlan all**. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the **vlan all** configuration. For example, in the following configuration:

```
protocols {
  mld-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee {
      interface ge-0/0/8.0 {
        static {
          group ff1e::1;
        }
      }
    }
  }
}
```

all VLANs, except **employee**, have a robust count of 8. Because **employee** has been individually configured, its **robust count** value is not determined by the value set under **vlan all**. Instead, its robust count is the default value of 2.

Enabling or Disabling MLD Snooping on VLANs

MLD snooping is not enabled on any VLAN by default. You must explicitly configure a VLAN or all VLANs for MLD snooping.

This topic describes how you can enable or disable MLD snooping on specific VLANs or on all VLANs on the switch.

- To enable MLD snooping on all VLANs:

```
[edit protocols mld-snooping]
user@switch# set vlan all
```

- To enable MLD snooping on a specific VLAN:

```
[edit protocols mld-snooping]
user@switch# set vlan vlan-name
```



NOTE: You cannot configure MLD snooping on a secondary VLAN.

For example, to enable MLD snooping on VLAN **education**:

```
[edit protocols mld-snooping]
user@switch# set vlan education
```

- To enable MLD snooping on all VLANs except a few VLANs:

1. Enable MLD snooping on all VLANs:

```
[edit protocols mld-snooping]
user@switch# set vlan all
```


2. Disable MLD snooping on individual VLANs:

```
[edit protocols mld-snooping]
user@switch# set vlan vlan-name disable
```

For example, to enable MLD snooping on all VLANs except **vlan100** and **vlan200**:

```
[edit protocols mld-snooping]
user@switch# set vlan all
```

```
[edit protocols mld-snooping]
user@switch# set vlan vlan100 disable
```

```
[edit protocols mld-snooping]
user@switch# set vlan vlan200 disable
```

You can also deactivate the MLD snooping protocol on the switch without changing the MLD snooping VLAN configurations:

```
[edit]
user@switch# deactivate protocols mld-snooping
```

Configuring the MLD Version

You can configure the version of MLD queries sent by a switch when MLD snooping is enabled. By default, the switch uses MLD version 1 (MLDv1). If you are using Protocol-Independent Multicast source-specific multicast (PIM-SSM), we recommend that you configure the switch to use MLDv2.

Typically, a switch passively monitors MLD messages sent between multicast routers and hosts and does not send MLD queries. The exception is when a switch detects that an interface has come up. When an interface comes up, the switch sends an immediate general membership query to all hosts on the interface. By doing so, the switch enables the multicast routers to learn group memberships more quickly than they would if they had to wait until the MLD querier sent its next general query.

The MLD version of the general query determines the MLD version of the host membership reports as follows:

- MLD version 1 (MLDv1) general query—Both MLDv1 and MLDv2 hosts respond with an MLDv1 membership report.
- MLDv2 general query—MLDv2 hosts respond with an MLDv2 membership report, while MLDv1 hosts are unable to respond to the query.

By default, the switch sends MLDv1 queries. This ensures compatibility with hosts and multicast routers that support MLDv1 only and cannot process MLDv2 reports. However, if your VLAN contains MLDv2 multicast routers and hosts and the routers are running PIM-SSM, we recommend that you configure MLD snooping for MLDv2. Doing so enables the routers to quickly learn which multicast sources the hosts on the interface want to receive traffic from.



NOTE: Configuring the MLD version does not limit the version of MLD messages that the switch can snoop. A switch can snoop both MLDv1 and MLDv2 messages regardless of the MLD version configured.

To configure the MLD version on a switch:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name version number
```

For example, to set the MLD version to version 2 for VLAN **marketing**:

```
[edit protocols]user@switch# set mld-snooping vlan marketing version 2
```

Enabling Immediate Leave

By default, when a switch with MLD snooping enabled receives an MLD leave report on a member interface, it waits for hosts on the interface to respond to MLD group-specific queries to determine whether there still are hosts on the interface interested in receiving the group multicast traffic. If the switch does not see any membership reports for the group within a set interval of time, it removes the interface's group membership from the multicast forwarding table and stops forwarding multicast traffic for the group to the interface.

You can decrease the leave latency created by this default behavior by enabling immediate leave on a VLAN.

When you enable immediate leave on a VLAN, host tracking is also enabled, allowing the switch to keep track of the hosts on a interface that have joined a multicast group. When the switch receives a leave report from the last member of the group, it immediately stops forwarding traffic to the interface and does not wait for the interface group membership to time out.

Immediate leave is supported for both MLD version 1 (MLDv1) and MLDv2. However, with MLDv1, we recommend that you configure immediate leave only when there is only one MLD host on an interface. In MLDv1, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports. This report-suppression feature means that the switch only knows about one interested host at any given time.

To enable immediate leave on a VLAN:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name immediate-leave
```

To enable immediate leave on all VLANs:

```
[edit protocols]user@switch# set mld-snooping vlan all immediate-leave
```

Configuring an Interface as a Multicast-Router Interface

When MLD snooping is enabled on a switch, the switch determines which interfaces face a multicast router by monitoring interfaces for MLD queries or Protocol Independent Multicast (PIM) updates. If the switch receives these messages on an interface, it adds the interface to its multicast forwarding table as a multicast-router interface.

In addition to dynamically learned interfaces, the multicast forwarding table can include interfaces that you explicitly configure to be multicast router interfaces. Unlike the table entries for dynamically learned interfaces, table entries for statically configured interfaces are not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure a static multicast-router interface include:

- You have an unusual network configuration that prevents MLD snooping from reliably learning about a multicast-router interface through monitoring MLD queries or PIM updates.
- Your implementation does not require an MLD querier.
- You have a stable topology and want to avoid the delay the dynamic learning process entails.



NOTE: If the interface you are configuring as a multicast-router interface is a trunk port, the interface becomes a multicast-router interface for all VLANs configured on the trunk port even if you have not explicitly configured it for all the VLANs. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast-router interface, even if the interface is configured as a multicast-router interface only for MLD snooping.

To configure an interface as a static multicast-router interface:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name interface interface-name multicast-router-interface
```

For example, to configure **ge-0/0/5.0** as a multicast-router interface for all VLANs on the switch:

```
[edit protocols]user@switch# set mld-snooping vlan all interface ge-0/0/5.0 multicast-router-interface
```

Configuring Static Group Membership on an Interface

To determine how to forward multicast packets, a switch with MLD snooping enabled maintains a multicast forwarding table containing a list of host interfaces that have interested listeners for a specific multicast group. The switch learns which host interfaces to add or delete from this table by examining MLD membership reports as they arrive on interfaces on which MLD snooping is enabled.

In addition to such dynamically learned interfaces, the multicast forwarding table can include interfaces that you statically configure to be members of multicast groups. When you configure a static group interface, the switch adds the interface to the forwarding table as a host interface for the group. Unlike an entry for a dynamically learned interface, a static interface entry is not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure static group membership on an interface include:

- You want to simulate an attached multicast receiver for testing purposes.
- The interface has receivers that cannot send MLD membership reports.
- You want the multicast traffic for a specific group to be immediately available to a receiver without any delay imposed by the dynamic join process.

You cannot configure multicast source addresses for a static group interface. The MLD version of a static group interface is always MLD version 1.



NOTE: The switch does not simulate MLD membership reports on behalf of a statically configured interface. Thus a multicast router might be unaware that the switch has an interface that is a member of the multicast group. You can configure a static group interface on the router to ensure that the switch receives the group multicast traffic.

To configure a host interface as a static member of a multicast group:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name interface interface-name static
group ip-address
```

For example, to configure interface **ge-0/0/11.0** in VLAN **ip-camera-vlan** as a static member of multicast group **ff1e::1**:

```
[edit protocols]user@switch# set mld-snooping vlan ip-camera-vlan interface ge-0/0/11.0 static
group ff1e::1
```

Changing the Timer and Counter Values

MLD uses various timers and counters to determine how often an MLD querier sends out membership queries and when group memberships time out. On Juniper Networks EX Series switches, the MLD and MLD snooping timers and counters default values are set to the values recommended in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*. These values work well for most multicast implementations.

There might be cases, however, where you might want to adjust the timer and counter values—for example, to reduce burstiness, to reduce leave latency, or to adjust for expected packet loss on a subnet. If you change a timer or counter value for the MLD querier on a VLAN, we recommend that you change the value for all multicast routers and switches on the VLAN so that all devices time out group memberships at approximately the same time.

The following timers and counters are configurable on a switch:

- **query-interval**—The length of time the MLD querier waits between sending general queries (the default is 125 seconds). You can change this interval to tune the number of MLD messages on the subnet; larger values cause general queries to be sent less often.

You cannot configure this value directly for MLD snooping. MLD snooping inherits the value from the MLD value configured on the switch, which is applied to all VLANs on the switch.

To configure the MLD **query-interval**:

```
[edit protocols]user@switch# set mld query-interval seconds
```

- **query-response-interval**—The maximum length of time the host can wait until it responds (the default is 10 seconds). You can change this interval to adjust the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

You cannot configure this value directly for MLD snooping. MLD snooping inherits the value from the MLD value configured on the switch, which is applied to all VLANs on the switch.

To configure the MLD **query-response-interval**:

```
[edit protocols]user@switch# set mld query-response-interval seconds
```

- **query-last-member-interval**—The length of time the MLD querier waits between sending group-specific membership queries (the default is 1 second). The MLD querier sends a group-specific query after receiving a leave report from a host. You can decrease this interval to reduce the amount of time it takes for multicast traffic to stop forwarding after the last member leaves a group.

You cannot configure this value directly for MLD snooping. MLD snooping inherits the value from the MLD value configured on the switch, which is applied to all VLANs on the switch.

To configure the MLD **query-last-member-interval**:

```
[edit protocols]user@switch# set mld query-last-member-interval seconds
```

- **robust-count**—The number of times the querier resends a general membership query or a group-specific membership query (the default is 2 times). You can increase this count to tune for higher expected packet loss.

For MLD snooping, you can configure **robust-count** for a specific VLAN. If a VLAN does not have **robust-count** configured, the **robust-count** value is inherited from the value configured for MLD.

To configure **robust-count** for MLD snooping on a VLAN:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name robust-count number
```

The values configured for **query-interval**, **query-response-interval**, and **robust-count** determine the multicast listener interval—the length of time the switch waits for a group membership report after a general query before removing a multicast group from its multicast forwarding table. The switch calculates the multicast listener interval by multiplying **query-interval** by **robust-count** and then adding **query-response-interval**:

(query-interval x robust-count) + query-response-interval = multicast listener interval

For example, the multicast listener interval is 260 seconds when the default settings for **query-interval**, **query-response-interval**, and **robust-count** are used:

$(125 \times 2) + 10 = 260$

You can display the time remaining in the multicast listener interval before a group times out by using the `show mld-snooping membership` command.

RELATED DOCUMENTATION

| [Configuring MLD | 59](#)

Configuring MLD Snooping on a Switch VLAN with ELS Support (CLI Procedure)

IN THIS SECTION

- [Enabling or Disabling MLD Snooping on VLANs | 198](#)
- [Configuring the MLD Version | 199](#)
- [Enabling Immediate Leave | 200](#)
- [Configuring an Interface as a Multicast-Router Interface | 200](#)
- [Configuring Static Group Membership on an Interface | 201](#)
- [Changing the Timer and Counter Values | 202](#)



NOTE: This task uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see ["Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\)" on page 188](#). For ELS details, see [Using the Enhanced Layer 2 Software CLI](#).

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on the VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

You can perform the following configurations for each VLAN:

- Selectively enable MLD snooping on specific VLANs.

- Specify the MLD version for the general query that the switch sends on an interface when the interface comes up.
- Enable immediate leave to reduce the length of time it takes the switch to stop forwarding multicast traffic when the last member host on the interface leaves the group.
- Configure an interface as a static multicast-router interface so that the switch does not need to dynamically learn that the interface is a multicast-router interface.
- Configure an interface as a static member of a multicast group so that the switch does not need to dynamically learn the interface's membership.
- Change the value for certain timers and counters to match the values configured on the multicast router serving as the MLD querier.

Enabling or Disabling MLD Snooping on VLANs

MLD snooping is not enabled on any VLAN by default. You must explicitly enable MLD snooping on specific interfaces.

- To enable MLD snooping on a specific VLAN:

```
[edit protocols mld-snooping]
user@switch# set vlan vlan-name
```



NOTE: You cannot enable MLD snooping on a secondary VLAN.

For example, to enable MLD snooping on VLAN education:

```
[edit protocols mld-snooping]
user@switch# set vlan education
```

- To disable MLD snooping on a specific VLAN:

```
[edit protocols mld-snooping]
user@switch# delete vlan vlan-name
```


You can also deactivate the MLD snooping protocol on the switch without changing the MLD snooping VLAN configurations:

```
[edit]
user@switch# deactivate protocols mld-snooping
```

Configuring the MLD Version

You can configure the version of MLD queries sent by a switch when MLD snooping is enabled. By default, the switch uses MLD version 1 (MLDv1). If you are using Protocol-Independent Multicast source-specific multicast (PIM-SSM), we recommend that you configure the switch to use MLDv2.

Typically, a switch passively monitors MLD messages sent between multicast routers and hosts and does not send MLD queries. The exception is when a switch detects that an interface has come up. When an interface comes up, the switch sends an immediate general membership query to all hosts on the interface. By doing so, the switch enables the multicast routers to learn group memberships more quickly than they would if they had to wait until the MLD querier sent its next general query.

The MLD version of the general query determines the MLD version of the host membership reports as follows:

- MLD version 1 (MLDv1) general query—Both MLDv1 and MLDv2 hosts respond with an MLDv1 membership report.
- MLDv2 general query—MLDv2 hosts respond with an MLDv2 membership report, while MLDv1 hosts are unable to respond to the query.

By default, the switch sends MLDv1 queries. This ensures compatibility with hosts and multicast routers that support MLDv1 only and cannot process MLDv2 reports. However, if your VLAN contains MLDv2 multicast routers and hosts and the routers are running PIM-SSM, we recommend that you configure MLD snooping for MLDv2. Doing so enables the routers to quickly learn which multicast sources the hosts on the interface want to receive traffic from.



NOTE: Configuring the MLD version does not limit the version of MLD messages that the switch can snoop. A switch can snoop both MLDv1 and MLDv2 messages regardless of the MLD version configured.

To configure the MLD version on an interface:

```
[edit protocols]user@switch# set mld interface interface-name version number
```

For example, to set the MLD version to version 2 on interface ge-0/0/2:

```
[edit protocols]user@switch# set mld interface ge-0/0/2 version 2
```

Enabling Immediate Leave

By default, when a switch with MLD snooping enabled receives an MLD leave report on a member interface, it waits for hosts on the interface to respond to MLD group-specific queries to determine whether there still are hosts on the interface interested in receiving the group multicast traffic. If the switch does not see any membership reports for the group within a set interval of time, it removes the interface's group membership from the multicast forwarding table and stops forwarding multicast traffic for the group to the interface.

You can decrease the leave latency created by this default behavior by enabling immediate leave on a VLAN.

When you enable immediate leave on a VLAN, host tracking is also enabled, allowing the switch to keep track of the hosts on a interface that have joined a multicast group. When the switch receives a leave report from the last member of the group, it immediately stops forwarding traffic to the interface and does not wait for the interface group membership to time out.

Immediate leave is supported for both MLD version 1 (MLDv1) and MLDv2. However, with MLDv1, we recommend that you configure immediate leave only when there is only one MLD host on an interface. In MLDv1, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports. This report-suppression feature means that the switch only knows about one interested host at any given time.

To enable immediate leave on a VLAN:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name immediate-leave
```

Configuring an Interface as a Multicast-Router Interface

When MLD snooping is enabled on a switch, the switch determines which interfaces face a multicast router by monitoring interfaces for MLD queries or Protocol Independent Multicast (PIM) updates. If the switch receives these messages on an interface, it adds the interface to its multicast forwarding table as a multicast-router interface.

In addition to dynamically learned interfaces, the multicast forwarding table can include interfaces that you explicitly configure to be multicast router interfaces. Unlike the table entries for dynamically learned interfaces, table entries for statically configured interfaces are not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure a static multicast-router interface include:

- You have an unusual network configuration that prevents MLD snooping from reliably learning about a multicast-router interface through monitoring MLD queries or PIM updates.
- Your implementation does not require an MLD querier.
- You have a stable topology and want to avoid the delay the dynamic learning process entails.

To configure an interface as a static multicast-router interface:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name interface interface-name multicast-router-interface
```

For example, to configure ge-0/0/5.0 as a multicast-router interface for VLAN employee:

```
[edit protocols]user@switch# set mld-snooping vlan employee interface ge-0/0/5.0 multicast-router-interface
```

Configuring Static Group Membership on an Interface

To determine how to forward multicast packets, a switch with MLD snooping enabled maintains a multicast forwarding table containing a list of host interfaces that have interested listeners for a specific multicast group. The switch learns which host interfaces to add or delete from this table by examining MLD membership reports as they arrive on interfaces on which MLD snooping is enabled.

In addition to such dynamically learned interfaces, the multicast forwarding table can include interfaces that you statically configure to be members of multicast groups. When you configure a static group interface, the switch adds the interface to the forwarding table as a host interface for the group. Unlike an entry for a dynamically learned interface, a static interface entry is not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure static group membership on an interface include:

- You want to simulate an attached multicast receiver for testing purposes.
- The interface has receivers that cannot send MLD membership reports.
- You want the multicast traffic for a specific group to be immediately available to a receiver without any delay imposed by the dynamic join process.

You cannot configure multicast source addresses for a static group interface. The MLD version of a static group interface is always MLD version 1.



NOTE: The switch does not simulate MLD membership reports on behalf of a statically configured interface. Thus a multicast router might be unaware that the switch has an interface that is a member of the multicast group. You can configure a static group interface on the router to ensure that the switch receives the group multicast traffic.

To configure a host interface as a static member of a multicast group:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name interface interface-name static
group ip-address
```

For example, to configure interface ge-0/0/11.0 in VLAN employee as a static member of multicast group ff1e::1:

```
[edit protocols]user@switch# set mld-snooping vlan ip-camera-vlan interface ge-0/0/11.0 static
group ff1e::1
```

Changing the Timer and Counter Values

MLD uses various timers and counters to determine how often an MLD querier sends out membership queries and when group memberships time out. On Juniper Networks switches, the MLD and MLD snooping timers and counters default values are set to the values recommended in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*. These values work well for most IPv6 multicast deployments.

There might be cases, however, where you might want to adjust the timer and counter values—for example, to reduce burstiness, to reduce leave latency, or to adjust for expected packet loss on a subnet. If you change a timer or counter value for the MLD querier on a VLAN, we recommend that you change the value for all multicast routers and switches on the VLAN so that all devices time out group memberships at approximately the same time.

The following timers and counters are configurable on a switch:

- `query-interval`—The length of time in seconds the MLD querier waits between sending general queries (the default is 125 seconds). You can change this interval to tune the number of MLD messages on the subnet; larger values cause general queries to be sent less often.

To configure the MLD query interval:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name query-interval seconds
```

- `query-response-interval`—The maximum length of time in seconds the host waits before it responds (the default is 10 seconds). You can change this interval to accommodate the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

To configure the MLD query response interval:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name query-response-interval seconds
```

- `query-last-member-interval`—The length of time the MLD querier waits between sending group-specific membership queries (the default is 1 second). The MLD querier sends a group-specific query after receiving a leave report from a host. You can decrease this interval to reduce the amount of time it takes for multicast traffic to stop forwarding after the last member leaves a group.

To configure the MLD query last member interval:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name query-last-member-interval seconds
```

- `robust-count`—The number of times the querier resends a general membership query or a group-specific membership query (the default is 2 times). You can increase this count to tune for higher anticipated packet loss.

For MLD snooping, you can configure `robust-count` for a specific VLAN. If a VLAN does not have `robust-count` configured, the value is inherited from the value configured for MLD.

To configure `robust-count` for MLD snooping on a VLAN:

```
[edit protocols]user@switch# set mld-snooping vlan vlan-name robust-count number
```

The values configured for `query-interval`, `query-response-interval`, and `robust-count` determine the multicast listener interval—the length of time the switch waits for a group membership report after a general query before removing a multicast group from its multicast forwarding table. The switch calculates the multicast listener interval by multiplying `query-interval` value by the `robust-count` value and then adding the `query-response-interval` to the product:

$$(\text{query-interval} \times \text{robust-count}) + \text{query-response-interval} = \text{multicast listener interval}$$

For example, the multicast listener interval is 260 seconds when the default settings for `query-interval`, `query-response-interval`, and `robust-count` are used:

$$(125 \times 2) + 10 = 260$$

To display the time remaining in the multicast listener interval before a group times out, use the `show mld-snooping membership` command.

RELATED DOCUMENTATION

[Example: Configuring MLD Snooping on Switches with ELS Support | 228](#)

[Configuring MLD | 59](#)

[Verifying MLD Snooping on Switches | 239](#)

Example: Configuring MLD Snooping on EX Series Switches

IN THIS SECTION

- [Requirements | 204](#)
- [Overview and Topology | 205](#)
- [Configuration | 207](#)
- [Verifying MLD Snooping Configuration | 208](#)

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure MLD snooping:

Requirements

This example uses the following software and hardware components:

- One EX Series switch
- Junos OS Release 12.1 or later

Before you configure MLD snooping, be sure you have:

- Configured the **vlan100** VLAN on the switch

- Assigned interfaces **ge-0/0/0**, **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/12** to **vlan100**
- Configured **ge-0/0/12** as a trunk interface.

See [Configuring VLANs for EX Series Switches](#).

Overview and Topology

IN THIS SECTION


- [Topology | 205](#)

In this example, interfaces **ge-0/0/0**, **ge-0/0/1**, and **ge-0/0/2** on the switch are in **vlan100** and are connected to hosts that are potential multicast receivers. Interface **ge-0/0/12**, a trunk interface also in **vlan100**, is connected to a multicast router. The router acts as the MLD querier and forwards multicast traffic for group **ff1e::2010** to the switch from a multicast source.

Topology

The example topology is illustrated in [Figure 26 on page 206](#).

Figure 26: Example MLD Snooping Topology



In this example topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group **ff1e::2010** from one of the hosts—for example, Host B. If MLD snooping is not enabled on **vlan100**, the switch floods the multicast traffic on all interfaces in **vlan100** (except for interface **ge-0/0/12**). If MLD snooping is enabled on **vlan100**, the switch monitors the MLD messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface **ge-0/0/1**.

This example shows how to enable MLD snooping on **vlan100**. It also shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast group has left the group. If immediate leave is not configured, the switch waits until the group-specific membership queries time out before it stops forwarding traffic.
- Configure **ge-0/0/12** as a static multicast-router interface. In this topology, **ge-0/0/12** always leads to the multicast router. By statically configuring **ge-0/0/12** as a multicast-router interface, you avoid any delay imposed by the switch having to learn that **ge-0/0/12** is a multicast-router interface.

Configuration

IN THIS SECTION

- [Procedure](#) | 207

To configure MLD snooping on a switch:

Procedure

CLI Quick Configuration

To quickly configure MLD snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols mld-snooping vlan vlan100

set protocols mld-snooping vlan vlan100 immediate-leave
set protocols mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Step-by-Step Procedure

To configure MLD snooping:

1. Enable MLD snooping on VLAN **vlan100**:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100
```

2. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 immediate-leave
```

3. Statically configure interface **ge-0/0/12** as a multicast-router interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show mld-snooping
vlan vlan100 {
  immediate-leave;
  interface ge-0/0/12.0 {
    multicast-router-interface;
  }
}
```

Verifying MLD Snooping Configuration

IN THIS SECTION

- [Verifying MLD Snooping Interface Membership on VLAN vlan100 | 208](#)

To verify that MLD snooping is enabled on the VLAN and the MLD snooping forwarding interfaces are correct, perform the following task:

Verifying MLD Snooping Interface Membership on VLAN vlan100

Purpose

Verify that MLD snooping is enabled on **vlan100** and that the multicast-router interface is statically configured:

Action

Show the group memberships maintained by MLD snooping for **vlan100**:

```
user@switch> show mld-snooping membership vlan vlan100 detail
VLAN: vlan100 Tag: 100 (Index: 8)
  Router interfaces:
    ge-0/0/12.0 static Uptime: 00:15:03
  Group: ff1e::2010
    ge-0/0/1.0 Timeout: 225 Flags: <V2-hosts>
    Last reporter: fe80::2020:1:1:3
```

Meaning

MLD snooping is running on **vlan100**, and interface **ge-0/0/12.0** is a statically configured multicast-router interface. Because the multicast group **ff1e::2010** is listed, at least one host in the VLAN is a current member of the multicast group and that host is on interface **ge-0/0/1.0**.

RELATED DOCUMENTATION

[Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)

[Verifying MLD Snooping on EX Series Switches \(CLI Procedure\) | 234](#)

[Understanding MLD Snooping | 177](#)

Example: Configuring MLD Snooping on SRX Series Devices

IN THIS SECTION

- [Requirements | 210](#)
- [Overview and Topology | 210](#)
- [Configuration | 211](#)
- [Verifying MLD Snooping Configuration | 215](#)

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, SRX Series Firewall examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the device then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure MLD snooping:

Requirements

This example uses the following software and hardware components:

- One SRX Series Firewall
- Junos OS Release 18.1R1

Before you configure MLD snooping, be sure you have:

- Configured the `vlan100` VLAN on the device
- Assigned interfaces `ge-0/0/0`, `ge-0/0/1`, `ge-0/0/2`, and `ge-0/0/3` to `vlan100`
- Configured `ge-0/0/3` as a trunk interface.

Overview and Topology

IN THIS SECTION

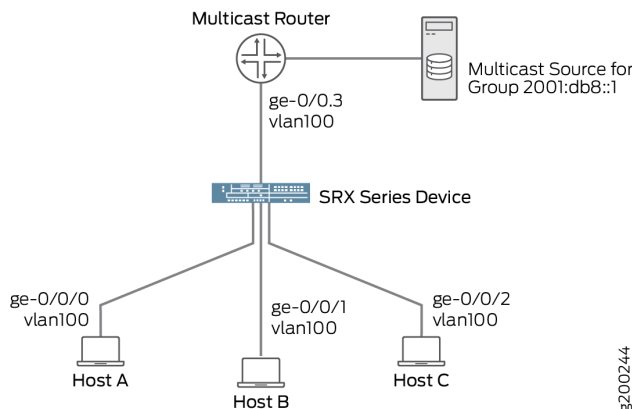
- [Topology | 210](#)

In this example, interfaces `ge-0/0/0`, `ge-0/0/1`, and `ge-0/0/2` on the device are in `vlan100` and are connected to hosts that are potential multicast receivers. Interface `ge-0/0/3`, a trunk interface also in `vlan100`, is connected to a multicast router. The router acts as the MLD querier and forwards multicast traffic for group `2001:db8::1` to the device from a multicast source.

Topology

The example topology is illustrated in [Figure 27 on page 211](#).

Figure 27: Example MLD Snooping Topology



In this example topology, the multicast router forwards multicast traffic to the device from the source when it receives a membership report for group `2001:db8::1` from one of the hosts—for example, Host B. If MLD snooping is not enabled on `vlan100`, then the device floods the multicast traffic on all interfaces in `vlan100` (except for interface `ge-0/0/3`). If MLD snooping is enabled on `vlan100`, the device monitors the MLD messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The device then forwards the multicast traffic only to interface `ge-0/0/1`.

This example shows how to enable MLD snooping on `vlan100`. It also shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the device stops forwarding multicast traffic on an interface when it detects that the last member of the multicast group has left the group. If immediate leave is not configured, the device waits until the group-specific membership queries time out before it stops forwarding traffic
- Configure `ge-0/0/3` as a static multicast-router interface. In this topology, `ge-0/0/3` always leads to the multicast router. By statically configuring `ge-0/0/3` as a multicast-router interface, you avoid any delay imposed by the device having to learn that `ge-0/0/3` is a multicast-router interface.

Configuration

IN THIS SECTION

- Procedure | [212](#)

To configure MLD snooping on a device:

Procedure

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.

```
set interfaces ge-0/0/0 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members vlan100
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members vlan100
set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members vlan100
set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members vlan100
set vlans vlan100 vlan-id 100
set routing-options nonstop-routing
set protocols mld-snooping vlan vlan100 query-interval 200
set protocols mld-snooping vlan vlan100 query-response-interval 0.4
set protocols mld-snooping vlan vlan100 query-last-member-interval 0.1
set protocols mld-snooping vlan vlan100 robust-count 4
set protocols mld-snooping vlan vlan100 immediate-leave
set protocols mld-snooping vlan vlan100 interface ge-0/0/1.0 host-only-interface
set protocols mld-snooping vlan vlan100 interface ge-0/0/0.0 group-limit 50
set protocols mld-snooping vlan vlan100 interface ge-0/0/2.0 static group 2001:db8::1
set protocols mld-snooping vlan vlan100 interface ge-0/0/3.0 multicast-router-interface
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

To configure MLD snooping:

1. Configure the access mode interfaces.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 family ethernet-switching interface-mode access
user@host# set ge-0/0/0 unit 0 family ethernet-switching vlan members vlan100
user@host# set ge-0/0/1 unit 0 family ethernet-switching interface-mode access
```

```

user@host# set ge-0/0/1 unit 0 family ethernet-switching vlan members vlan100
user@host# set ge-0/0/2 unit 0 family ethernet-switching interface-mode access
user@host# set ge-0/0/2 unit 0 family ethernet-switching vlan members vlan100

```

2. Configure the trunk mode interface.

```

[edit interfaces]
user@host# set ge-0/0/3 unit 0 family ethernet-switching interface-mode trunk
user@host# set ge-0/0/3 unit 0 family ethernet-switching vlan members vlan100

```

3. Configure the VLAN.

```

[edit vlans vlan100]
user@host# set vlans v100 vlan-id 100

```

4. Configure nonstop routing

```

[edit]
user@host# set routing-options nonstop-routing

```

5. Configure the limit for the number of multicast groups allowed on the ge-0/0/1.0 interface to 50.

```

[edit vlans vlan100]
user@host# set protocols mld-snooping vlan vlan100 interface ge-0/0/0.0 group-limit 50

```

6. Configure the device to immediately remove a group membership from an interface when it receives a leave message from that interface without waiting for any other MLD messages to be exchanged.

```

[edit vlans vlan100]
user@host# set protocols mld-snooping vlan vlan100 immediate-leave

```

7. Statically configure interface ge-0/0/2.0 as a multicast-router interface.

```
[edit vlans vlan100]
user@host# set protocols mld-snooping vlan vlan100 interface ge-0/0/2.0 static group
2001:db8::1
```

8. Configure an interface to be an exclusively router-facing interface (to receive multicast traffic).

```
[edit vlans vlan100]
user@host# set protocols mld-snooping vlan vlan100 interface ge-0/0/3.0 multicast-router-
interface
```

9. Configure an interface to be an exclusively host-facing interface (to drop MLD query messages).

```
[edit vlans vlan100]
user@host# set protocols mld-snooping vlan vlan100 interface ge-0/0/1.0 host-only-interface
```

10. Configure the IGMP message intervals and robustness count.

```
[edit vlans vlan100]
uer@host# set protocols mld-snooping vlan v100 query-interval 200
uer@host# set protocols mld-snooping vlan v100 query-response-interval 0.4
uer@host# set protocols mld-snooping vlan v100 query-last-member-interval 0.1
uer@host# set protocols mld-snooping vlan v1 robust-count 4
```

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

```
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the `show protocols mld-snooping` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show protocols mld-snooping
```



```
vlan vlan100 {
  query-interval 200;
  query-response-interval 0.4;
  query-last-member-interval 0.1;
  robust-count 4;
  immediate-leave;
  interface ge-0/0/1.0 {
    host-only-interface;
  }
  interface ge-0/0/0.0 {
    group-limit 50;
  }
  interface ge-0/0/2.0 {
    static {
      group 2001:db8::1;
    }
  }
  interface ge-0/0/3.0 {
    multicast-router-interface;
  }
}
```

Verifying MLD Snooping Configuration

IN THIS SECTION

- [Verifying MLD Snooping Interface Membership on VLAN vlan100 | 215](#)

To verify that MLD snooping is enabled on the VLAN and the MLD snooping forwarding interfaces are correct, perform the following task:

Verifying MLD Snooping Interface Membership on VLAN vlan100

Purpose

Verify that MLD snooping is enabled on `vlan100` and that the multicast-router interface is statically configured:

Action

From operational mode, enter the `show mld snooping membership` command.

```
user@host> show mld snooping membership
Instance: default-switch

Vlan: vlan100

Learning-Domain: default
Interface: ge-0/0/0.0, Groups: 0
Interface: ge-0/0/1.0, Groups: 0
Interface: ge-0/0/2.0, Groups: 1
  Group: 2001:db8::1
    Group mode: Exclude
    Source: ::
    Last reported by: Local
    Group timeout:      0 Type: Static
```

Meaning

MLD snooping is running on `vlan100`, and interface `ge-0/0/3.0` is a statically configured multicast-router interface. Because the multicast group `2001:db8::1` is listed, at least one host in the VLAN is a current member of the multicast group and that host is on interface `ge-0/0/1.0`.

RELATED DOCUMENTATION

mld-snooping

[Understanding MLD Snooping | 177](#)

Configuring MLD Snooping Tracing Operations on EX Series Switches (CLI Procedure)

IN THIS SECTION

- [Configuring Tracing Operations | 218](#)
- [Viewing, Stopping, and Restarting Tracing Operations | 219](#)

By enabling tracing operations for MLD snooping, you can record detailed messages about the operation of the protocol, such as the various types of protocol packets sent and received. [Table 9 on page 217](#) describes the tracing operations you can enable and the flags used to specify them in the tracing configuration.

Table 9: Supported Tracing Operations for MLD Snooping

Tracing Operation	Flag
Trace all (equivalent of including all flags).	all
Trace general MLD snooping protocol events.	general
Trace communication over routing socket events.	krt
Trace leave reports.	leave
Trace next-hop-related events.	nexthop
Trace normal MLD snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced.	normal
Trace all MLD packets.	packets
Trace policy processing.	policy

Table 9: Supported Tracing Operations for MLD Snooping (Continued)

Tracing Operation	Flag
Trace MLD membership query messages.	query
Trace membership reports	report
Trace routing information.	route
Trace state transitions.	state
Trace routing protocol task processing.	task
Trace timer processing.	timer
Trace VLAN-related events.	vlan

Configuring Tracing Operations

To configure tracing operations for MLD snooping:

1. Configure the filename for the trace file:

```
[edit protocols mld-snooping ]user@switch# set traceoptions file filename
```

For example:

```
[edit protocols mld-snooping ]user@switch# set traceoptions file mld-snoop-trace
```

2. (Optional) Configure the maximum number of trace files and size of the trace files:

```
[edit protocols mld-snooping ]user@switch # set file files number size size
```

For example:

```
[edit protocols mld-snooping ]user@switch # set traceoptions file files 5 size 1m
```

causes the contents of the trace file to be emptied and archived in a .gz file when the file reaches 1 MB. Four archive files are maintained, the contents of which are rotated whenever the current active trace file is archived.

If you omit this step, the maximum number of trace files defaults to 10, with the maximum file size defaulting to 128 K.

3. Specify one of the tracing flags shown in [Table 9 on page 217](#):

```
[edit protocols mld-snooping ]user@switch # set traceoptions flag flagname
```

For example, to perform trace operations on VLAN-related events and MLD query messages:

```
[edit protocols mld-snooping ]user@switch# set traceoptions flag vlan
```

```
[edit protocols mld-snooping ]user@switch# set traceoptions flag query
```

Viewing, Stopping, and Restarting Tracing Operations

When you commit the configuration, tracing operations begin. You can view the trace file in the `/var/log` directory. For example:

```
user@switch> file show /var/log/mld-snoop-trace
```

You can stop and restart tracing operations by deactivating and reactivating the configuration:

```
[edit] user@switch# deactivate protocols mld-snooping traceoptions
```

```
[edit] user@switch# activate protocols mld-snooping traceoptions
```

RELATED DOCUMENTATION

[Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)

[Tracing and Logging Junos OS Operations](#)

Configuring MLD Snooping Tracing Operations on EX Series Switch VLANs (CLI Procedure)

IN THIS SECTION

- [Configuring Tracing Operations | 221](#)
- [Viewing, Stopping, and Restarting Tracing Operations | 222](#)

By enabling tracing operations for MLD snooping, you can record detailed messages about the operation of the protocol, such as the various types of protocol packets sent and received. [Table 10 on page 220](#) describes the tracing operations you can enable and the flags used to specify them in the tracing configuration.

Table 10: Supported Tracing Operations for MLD Snooping

Tracing Operation	Flag
Trace all (equivalent of including all flags).	all
Trace client notifications.	client-notification
Trace general MLD snooping protocol events.	general
Trace group operations.	group
Trace host notifications.	host-notification

Table 10: Supported Tracing Operations for MLD Snooping (Continued)

Tracing Operation	Flag
Trace leave reports.	leave
Trace normal MLD snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced.	normal
Trace all MLD packets.	packets
Trace policy processing.	policy
Trace MLD membership query messages.	query
Trace membership reports.	report
Trace routing information.	route
Trace state transitions.	state
Trace routing protocol task processing.	task
Trace timer processing.	timer

Configuring Tracing Operations

To configure tracing operations for MLD snooping:

1. Configure the filename for the trace file:

```
[edit protocols mld-snooping ]user@switch# set vlan vlan-name traceoptions file filename
```

For example:

```
[edit protocols mld-snooping ]user@switch# set vlan vlan100 traceoptions file mld-snoop-trace
```

2. (Optional) Configure the maximum number of trace files and size of the trace files:

```
[edit protocols mld-snooping ]user@switch # set vlan vlan-name traceoptions file files number
size size
```

For example:

```
[edit protocols mld-snooping ]user@switch # set vlan vlan100 traceoptions file files 5 size 1m
```

causes the contents of the trace file to be emptied and archived in a .gz file when the file reaches 1 MB. Four archive files are maintained, the contents of which are rotated whenever the current active trace file is archived.

If you omit this step, the maximum number of trace files defaults to 10, and the maximum file size to 128 KB.

3. Specify one of the tracing flags shown in [Table 10 on page 220](#):

```
[edit protocols mld-snooping ]user@switch # set vlan vlan-name traceoptions flag flagname
```

For example, to perform trace operations on VLAN-related events and on MLD query messages:

```
[edit protocols mld-snooping ]user@switch# set vlan vlan100 traceoptions flag vlan
```

```
[edit protocols mld-snooping ]user@switch# set vlan vlan100 traceoptions flag query
```

Viewing, Stopping, and Restarting Tracing Operations

When you commit the configuration, tracing operations begin. You can view the trace file in the `/var/log` directory. For example:

```
user@switch> file show /var/log/mld-snoop-trace
```


You can stop and restart tracing operations by deactivating and reactivating the configuration:

```
[edit] user@switch# deactivate protocols mld-snooping traceoptions
```

```
[edit] user@switch# activate protocols mld-snooping traceoptions
```

RELATED DOCUMENTATION

[Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)

[Configuring MLD Snooping on a Switch VLAN with ELS Support \(CLI Procedure\) | 197](#)

[Tracing and Logging Junos OS Operations](#)

Example: Configuring MLD Snooping on EX Series Switches

IN THIS SECTION

- [Requirements | 223](#)
- [Overview and Topology | 224](#)
- [Configuration | 226](#)
- [Verifying MLD Snooping Configuration | 227](#)

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure MLD snooping:

Requirements

This example uses the following software and hardware components:

- One EX Series switch
- Junos OS Release 12.1 or later

Before you configure MLD snooping, be sure you have:

- Configured the **vlan100** VLAN on the switch
- Assigned interfaces **ge-0/0/0**, **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/12** to **vlan100**
- Configured **ge-0/0/12** as a trunk interface.

See [Configuring VLANs for EX Series Switches](#).

Overview and Topology

IN THIS SECTION


- [Topology | 224](#)

In this example, interfaces **ge-0/0/0**, **ge-0/0/1**, and **ge-0/0/2** on the switch are in **vlan100** and are connected to hosts that are potential multicast receivers. Interface **ge-0/0/12**, a trunk interface also in **vlan100**, is connected to a multicast router. The router acts as the MLD querier and forwards multicast traffic for group **ff1e::2010** to the switch from a multicast source.

Topology

The example topology is illustrated in [Figure 26 on page 206](#).

Figure 28: Example MLD Snooping Topology



In this example topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group **ff1e::2010** from one of the hosts—for example, Host B. If MLD snooping is not enabled on **vlan100**, the switch floods the multicast traffic on all interfaces in **vlan100** (except for interface **ge-0/0/12**). If MLD snooping is enabled on **vlan100**, the switch monitors the MLD messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface **ge-0/0/1**.

This example shows how to enable MLD snooping on **vlan100**. It also shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast group has left the group. If immediate leave is not configured, the switch waits until the group-specific membership queries time out before it stops forwarding traffic.
- Configure **ge-0/0/12** as a static multicast-router interface. In this topology, **ge-0/0/12** always leads to the multicast router. By statically configuring **ge-0/0/12** as a multicast-router interface, you avoid any delay imposed by the switch having to learn that **ge-0/0/12** is a multicast-router interface.

Configuration

IN THIS SECTION

- [Procedure](#) | 226

To configure MLD snooping on a switch:

Procedure

CLI Quick Configuration

To quickly configure MLD snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols mld-snooping vlan vlan100

set protocols mld-snooping vlan vlan100 immediate-leave
set protocols mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Step-by-Step Procedure

To configure MLD snooping:

1. Enable MLD snooping on VLAN **vlan100**:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100
```

2. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 immediate-leave
```

3. Statically configure interface **ge-0/0/12** as a multicast-router interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show mld-snooping
vlan vlan100 {
    immediate-leave;
    interface ge-0/0/12.0 {
        multicast-router-interface;
    }
}
```

Verifying MLD Snooping Configuration

IN THIS SECTION

- [Verifying MLD Snooping Interface Membership on VLAN vlan100 | 227](#)

To verify that MLD snooping is enabled on the VLAN and the MLD snooping forwarding interfaces are correct, perform the following task:

Verifying MLD Snooping Interface Membership on VLAN vlan100

Purpose

Verify that MLD snooping is enabled on **vlan100** and that the multicast-router interface is statically configured:

Action

Show the group memberships maintained by MLD snooping for **vlan100**:

```
user@switch> show mld-snooping membership vlan vlan100 detail
VLAN: vlan100 Tag: 100 (Index: 8)
  Router interfaces:
    ge-0/0/12.0 static Uptime: 00:15:03
  Group: ff1e::2010
    ge-0/0/1.0 Timeout: 225 Flags: <V2-hosts>
    Last reporter: fe80::2020:1:1:3
```

Meaning

MLD snooping is running on **vlan100**, and interface **ge-0/0/12.0** is a statically configured multicast-router interface. Because the multicast group **ff1e::2010** is listed, at least one host in the VLAN is a current member of the multicast group and that host is on interface **ge-0/0/1.0**.

RELATED DOCUMENTATION

[Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)

[Verifying MLD Snooping on EX Series Switches \(CLI Procedure\) | 234](#)

[Understanding MLD Snooping | 177](#)

Example: Configuring MLD Snooping on Switches with ELS Support

IN THIS SECTION

- [Requirements | 229](#)
- [Overview and Topology | 229](#)
- [Configuration | 231](#)
- [Verifying MLD Snooping Configuration | 232](#)



NOTE: This example uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [Using the Enhanced Layer 2 Software CLI](#).

You can enable MLD snooping on a VLAN to constrain the flooding of IPv6 multicast traffic on a VLAN. When MLD snooping is enabled, a switch examines MLD messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. On the basis of what it learns, the switch then forwards IPv6 multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces.

This example describes how to configure MLD snooping:

Requirements

This example uses the following software and hardware components:

- One switch running Junos OS with ELS
- Junos OS Release 13.3 or later for EX Series switches or Junos OS Release 15.1X53-D10 or later for QFX10000 switches

Before you configure MLD snooping, be sure you have:

- Configured the vlan 100 VLAN on the switch.
- Assigned interfaces ge-0/0/0, ge-0/0/1, ge-0/0/2, and ge-0/0/12 to vlan100.
- Configured ge-0/0/12 as a trunk interface.

See [Configuring VLANs for EX Series Switches](#) or [Configuring VLANs on Switches with Enhanced Layer 2 Support](#).

Overview and Topology

IN THIS SECTION


- [Topology | 230](#)

In this example, interfaces ge-0/0/0, ge-0/0/1, and ge-0/0/2 on the switch are in vlan100 and are connected to hosts that are potential multicast receivers. Interface ge-0/0/12, a trunk interface also in vlan100, is connected to a multicast router. The router acts as the MLD querier and forwards multicast traffic for group ff1e::2010 to the switch from a multicast source.

Topology

The topology for this example is illustrated in [Figure 29 on page 230](#).

Figure 29: MLD Snooping Topology Example



In this sample topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group ff1e::2010 from one of the hosts—for example, Host B. If MLD snooping is not enabled on vlan100, the switch floods the multicast traffic on all interfaces in vlan100 (except for interface ge-0/0/12). If MLD snooping is enabled on vlan100, the switch monitors the MLD messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface ge-0/0/1.

This example shows how to enable MLD snooping on vlan100. It also shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast

group has left the group. If immediate leave is not configured, the switch waits until the group-specific membership queries time out before it stops forwarding traffic.

- Configure ge-0/0/12 as a static multicast-router interface. In this topology, ge-0/0/12 always leads to the multicast router. By statically configuring ge-0/0/12 as a multicast-router interface, you avoid any delay imposed by the switch having to learn that ge-0/0/12 is a multicast-router interface.

Configuration

IN THIS SECTION

- [Procedure | 231](#)

To configure MLD snooping on a switch:

Procedure

CLI Quick Configuration

To quickly configure MLD snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols mld-snooping vlan vlan100

set protocols mld-snooping vlan vlan100 immediate-leave
set protocols mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Step-by-Step Procedure

To configure MLD snooping:

1. Enable MLD snooping on the VLAN vlan100:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100
```

2. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 immediate-leave
```

3. Statically configure interface ge-0/0/12 as a multicast-router interface:

```
[edit protocols]
user@switch# set mld-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

Results

Check the results of the configuration:

```
[edit protocols]
user@switch# show mld-snooping
vlan vlan100 {
    immediate-leave;
    interface ge-0/0/12.0 {
        multicast-router-interface;
    }
}
```

Verifying MLD Snooping Configuration

IN THIS SECTION

- [Verifying MLD Snooping Interface Membership on VLAN vlan100 | 233](#)

To verify that MLD snooping is enabled on the VLAN and the MLD snooping forwarding interfaces are correct, perform the following task:

Verifying MLD Snooping Interface Membership on VLAN vlan100

Purpose

Verify that MLD snooping is enabled on the VLAN vlan 100 and that the multicast-router interface is statically configured:

Action

Show the MLD snooping information for ge-0/0/12.0:

```
user@switch> show mld snooping interface
Instance: default-switch

Vlan: vlan100

Learning-Domain: default
Interface: ge-0/0/12.0
  State:          Up Groups:      3
  Immediate leave: On
  Router interface: yes

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

Meaning

MLD snooping is running on vlan100, and interface ge-0/0/12.0 is a statically configured multicast-router interface. Immediate leave is enabled on the interface.

RELATED DOCUMENTATION

[Configuring MLD Snooping on a Switch VLAN with ELS Support \(CLI Procedure\) | 197](#)

[Verifying MLD Snooping on Switches | 239](#)

[Understanding MLD Snooping | 177](#)

Verifying MLD Snooping on EX Series Switches (CLI Procedure)

IN THIS SECTION

- [Verifying MLD Snooping Memberships | 234](#)
- [Verifying MLD Snooping VLANs | 235](#)
- [Viewing MLD Snooping Statistics | 236](#)
- [Viewing MLD Snooping Routing Information | 237](#)

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs on a switch. This topic describes how to verify MLD snooping operation on the switch.

Verifying MLD Snooping Memberships

IN THIS SECTION

- [Purpose | 234](#)
- [Action | 234](#)
- [Meaning | 235](#)

Purpose

Determine group memberships, multicast-router interfaces, host MLD versions, and the current values of timeout counters.

Action

Enter the following command:

```
user@switch> show mld snooping membership detail
VLAN: mld-vlan Tag: 100 (Index: 3)
  Router interfaces:
    ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
  Group: ff1e::2010
```

```

ge-1/0/30.0 Timeout: 180 Flags: <V2-hosts>
Last reporter: fe80::2020:1:1:3
  Include source: 2020:1:1:1::2
  Include source: 2020:1:1:1::5

```

Meaning

The switch has multicast membership information for one VLAN on the switch, **mld-vlan**. MLD snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them. The following information is provided:

- Information on the multicast-router interfaces for the VLAN—in this case, **ge-1/0/0.0**. The multicast-router interface has been learned by MLD snooping, as indicated by **dynamic**. The **timeout** value shows how many seconds from now the interface will be removed from the multicast forwarding table if the switch does not receive MLD queries or Protocol Independent Multicast (PIM) updates on the interface.
- Information about the group memberships for the VLAN:
 - Currently, the VLAN has membership in only one multicast group, **ff1e::2010**.
 - The host or hosts that have reported membership in the group are on interface **ge-1/0/30.0**. The interface group membership will time out in 180 seconds if no hosts respond to membership queries during this interval. The flags field shows the lowest version of MLD used by a host that is currently a member of the group, which in this case is MLD version 2 (MLDv2).
 - The last host that reported membership in the group has address **fe80::2020:1:1:3**.
 - Because interface has MLDv2 hosts on it, the source addresses from which the MLDv2 hosts want to receive group multicast traffic are shown (addresses **2020:1:1:1::2** and **2020:1:1:1::5**). The **timeout** value for the interface group membership is derived from the largest timeout value for all sources addresses for the group.

Verifying MLD Snooping VLANs

IN THIS SECTION

- Purpose | 236
- Action | 236
- Meaning | 236

Purpose

Verify that MLD snooping is enabled on a VLAN and display MLD snooping information for each VLAN on which MLD snooping is enabled.

Action

Enter the following command:

```
user@switch> show mld-snooping vlans detail
VLAN: v10, Tag: 10
  Interface: ge-1/0/0.0, tagged, Groups: 0, Router
  Interface: ge-1/0/30.0, untagged, Groups: 1
  Interface: ge-12/0/30.0, untagged, Groups: 0
VLAN: v20, Tag: 20
  Interface: ge-1/0/0.0, tagged, Groups: 0, Router
  Interface: ge-1/0/31.0, untagged, Groups: 0
  Interface: ge-12/0/31.0, untagged, Groups: 1
```

Meaning

MLD snooping is configured on two VLANs on the switch: **v10** and **v20**. Each interface in each VLAN is listed and the following information is provided:

- Whether the interface is a trunk (**tagged**) or access (**untagged**) interface.
- How many multicast groups the interface belongs to.
- Whether the interface is a multicast-router interface (**Router**).

Viewing MLD Snooping Statistics

IN THIS SECTION

- [Purpose | 237](#)
- [Action | 237](#)
- [Meaning | 237](#)

Purpose

Display MLD snooping statistics, such as number of MLD queries, reports, and leaves received and how many of these MLD messages contained errors.

Action

Enter the following command:

```
user@switch> show mld snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0
```

MLD Type	Received	Transmitted	Recv Errors
Queries:	74295	0	0
Reports:	18148423	0	16333523
Leaves:	0	0	0
Other:	0	0	0

Meaning

The output shows how many MLD messages of each type—**Queries**, **Reports**, **Leaves**—the switch received or transmitted on interfaces on which MLD snooping is enabled. For each message type, it also shows the number of MLD packets the switch received that had errors—for example, packets that do not conform to the MLDv1 or MLDv2 standards. If the **Recv Errors** count increases, verify that the hosts are compliant with MLDv1 or MLDv2 standards. If the switch is unable to recognize the MLD message type for a packet, it counts the packet under **Receive unknown**.

Viewing MLD Snooping Routing Information

IN THIS SECTION

- [Purpose | 238](#)
- [Action | 238](#)
- [Meaning | 238](#)

Purpose

Display the next-hop information maintained in the multicast forwarding table.

Action

Enter the following command:

```
user@switch> show mld-snooping route detail
VLAN          Group          Next-hop
mld-vlan      ::0000:2010    1323
              Interfaces: ge-1/0/30.0, ge-1/0/33.0
VLAN          Group          Next-hop
mld-vlan      ff00::         1317
              Interfaces: ge-1/0/0.0, ge-1/0/33.0
VLAN          Group          Next-hop
mld-vlan      ::0000:0000    1317
              Interfaces: ge-1/0/0.0
VLAN          Group          Next-hop
mld-vlan1     ::0000:2010    1324
              Interfaces: ge-12/0/31.0
VLAN          Group          Next-hop
mld-vlan1     ff00::         1318
              Interfaces: ae200.0
VLAN          Group          Next-hop
mld-vlan1     ::0000:0000    1318
              Interfaces: ae200.0
```

Meaning

The output shows the next-hop interfaces for a given multicast group on a VLAN. Only the last 32 bits of the group address are shown because the switch uses only these bits in determining multicast routes. For example, route **::0000:2010** on **mld-vlan** has next-hop interfaces **ge-1/0/30.0** and **ge-1/0/33.0**.

RELATED DOCUMENTATION

clear mld snooping membership

clear mld snooping statistics

[Example: Configuring MLD Snooping on EX Series Switches | 204](#)

[Configuring MLD Snooping on an EX Series Switch VLAN \(CLI Procedure\) | 188](#)

Verifying MLD Snooping on Switches

IN THIS SECTION

- [Verifying MLD Snooping Memberships | 239](#)
- [Verifying MLD Snooping Interfaces | 240](#)
- [Viewing MLD Snooping Statistics | 242](#)
- [Viewing MLD Snooping Routing Information | 243](#)



NOTE: This topic uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see ["Verifying MLD Snooping on EX Series Switches \(CLI Procedure\)" on page 234](#). For ELS details, see [Using the Enhanced Layer 2 Software CLI](#).

Multicast Listener Discovery (MLD) snooping constrains the flooding of IPv6 multicast traffic on VLANs. This topic describes how to verify MLD snooping operation on a VLAN.

Verifying MLD Snooping Memberships

IN THIS SECTION

- [Purpose | 239](#)
- [Action | 240](#)
- [Meaning | 240](#)

Purpose

Verify that MLD snooping is enabled on a VLAN and determine group memberships.

Action

Enter the following command:

```
user@switch> show mld snooping membership detail
Instance: default-switch

Vlan: v1

Learning-Domain: default
Interface: ge-0/0/1.0, Groups: 1
  Group: ff05::1
    Group mode: Exclude
    Source: ::
    Last reported by: fe80::
    Group timeout: 259 Type: Dynamic
Interface: ge-0/0/2.0, Groups: 0
```

Meaning

The switch has multicast membership information for one VLAN on the switch, v1. MLD snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them.

- The following information is provided about the group memberships for the VLAN:
 - Currently, the VLAN has membership in only one multicast group, ff05::1.
 - The host or hosts that have reported membership in the group are on interface ge-0/0/1.0.
 - The last host that reported membership in the group has address fe80::.
 - The interface group membership will time out in 259 seconds if no hosts respond to membership queries during this interval.
 - The group membership has been learned by MLD snooping, as indicated by *Dynamic*.

Verifying MLD Snooping Interfaces

IN THIS SECTION

 Purpose | 241

- Action | 241
- Meaning | 241

Purpose

Display MLD snooping information for each interface on which MLD snooping is enabled.

Action

Enter the following command:

```
user@switch>show mld snooping interface
Instance: default-switch

Vlan: v100

Learning-Domain: default
Interface: ge-0/0/1.0
  State:          Up Groups:      1
  Immediate leave: Off
  Router interface: no
Interface: ge-0/0/2.0
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: no

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

Meaning

MLD snooping is configured on one VLAN on the switch, v100. Each interface in each VLAN is listed and the following information is provided:

- How many multicast groups the interface belongs to.

- Whether immediate leave has been configured for the interface.
- Whether the interface is a multicast-router interface.

The output also shows the configured parameters for the MLD querier.

Viewing MLD Snooping Statistics

IN THIS SECTION

- Purpose | 242
- Action | 242
- Meaning | 243

Purpose

Display MLD snooping statistics, such as number of MLD queries, reports, and leaves received and how many of these MLD messages contained errors.

Action

Enter the following command:

```

user@switch>show mld snooping statistics
Vlan: v1
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0            4      0
Listener Report (v1)     447          0      0
Listener Done (v1/v2)    0            0      0
Listener Report (v2)     0            0      0
Other Unknown types      0            0      0

Vlan: v2
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0            4      0
Listener Report (v1)     154          0      0
Listener Done (v1/v2)    0            0      0
Listener Report (v2)     0            0      0
Other Unknown types      0            0      0

```

```

Instance: default-switch
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0            8     0
Listener Report (v1)     601          0     0
Listener Done (v1/v2)    0            0     0
Listener Report (v2)     0            0     0
Other Unknown types      0            0     0

MLD Global Statistics
Bad Length              0
Bad Checksum            0
Bad Receive If         0
Rx non-local            0
Timed out               0

```

Meaning

The output shows how many MLD messages of each type—Queries, Done, Report—the switch received or transmitted on interfaces on which MLD snooping is enabled. For each message type, it also shows the number of MLD packets the switch received that had errors—for example, packets that do not conform to the MLDv1 or MLDv2 standards. If the Rx errors count increases, verify that the hosts are compliant with MLDv1 or MLDv2 standards. If the switch is unable to recognize the MLD message type for a packet, it counts the packet under Other Unknown types.

Viewing MLD Snooping Routing Information

IN THIS SECTION

- [Purpose | 243](#)
- [Action | 244](#)
- [Meaning | 244](#)

Purpose

Display the next-hop information maintained in the multicast snooping forwarding table.

Action

Enter the following command:

```
user@switch>show multicast snooping route
Nexthop Bulking: OFF

          Family: INET6

Group: ff00::/8
  Source: ::/128
  Vlan: v1

Group: ff02::1/128
  Source: ::/128
  Vlan: v1
  Downstream interface list:
  ge-1/0/16.0

Group: ff05::1/128
  Source: ::/128
  Vlan: v1
  Downstream interface list:
  ge-1/0/16.0

Group: ff06::1/128
  Source: ::/128
  Vlan: v1
  Downstream interface list:
  ge-1/0/16.0
```

Meaning

The output shows the next-hop interfaces for a given multicast group on a VLAN. For example, route ff02::1/128 on VLAN v1 has the next-hop interface ge-1/0/16.0.

RELATED DOCUMENTATION

[Example: Configuring MLD Snooping on Switches with ELS Support | 228](#)

[Configuring MLD Snooping on a Switch VLAN with ELS Support \(CLI Procedure\) | 197](#)

Configuring Multicast VLAN Registration

IN THIS CHAPTER

- [Understanding Multicast VLAN Registration | 245](#)
- [Configuring Multicast VLAN Registration on EX Series Switches | 255](#)
- [Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS | 267](#)

Understanding Multicast VLAN Registration

IN THIS SECTION

- [Benefits of Multicast VLAN Registration | 246](#)
- [How MVR Works | 246](#)
- [Recommended MVR Configurations in the Access Layer on ELS Switches | 250](#)

Multicast VLAN registration (MVR) enables more efficient distribution of IPTV multicast streams across an Ethernet ring-based Layer 2 network.

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to each requesting VLAN.

When you configure MVR, you create a *multicast VLAN* (MVLAN) that becomes the only VLAN over which IPTV multicast traffic flows throughout the Layer 2 network. Devices with MVR enabled selectively forward IPTV multicast traffic from interfaces on the MVLAN (source interfaces) to hosts that are connected to interfaces that are not part of the MVLAN that you designate as *MVR receiver ports*. MVR receiver ports can receive traffic from a port on the MVLAN but cannot send traffic onto the MVLAN, and those ports remain in their own VLANs for bandwidth and security reasons.

Benefits of Multicast VLAN Registration

- Reduces the bandwidth required to distribute IPTV multicast streams by eliminating duplication of multicast streams from the same source to interested receivers on different VLANs.

How MVR Works

MVR operates similarly to and in conjunction with Internet Group Management Protocol (IGMP) snooping. Both MVR and IGMP snooping monitor IGMP join and leave messages and build forwarding tables based on the media access control (MAC) addresses of the hosts sending those IGMP messages. Whereas IGMP snooping operates within a given VLAN to regulate multicast traffic, MVR can operate with hosts on different VLANs in a Layer 2 network to selectively deliver IPTV multicast traffic to any requesting hosts. This reduces the bandwidth needed to forward the traffic.



NOTE: MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

MVR Basics

MVR is not enabled by default on devices that support MVR. You explicitly configure an MVLAN and assign a range of multicast group addresses to it. That VLAN carries MVLAN traffic for the configured multicast groups. You then configure other VLANs to be MVR receiver VLANs that receive multicast streams from the MVLAN. When MVR is configured on a device, the device receives only one copy of each MVR multicast stream, and then replicates the stream only to the hosts that want to receive it, while forwarding all other types of multicast traffic without modification.

You can configure multiple MVLANS on a device, but they must have disjoint multicast group subnets. An MVR receiver VLAN can be associated with more than one MVLAN on the device.

MVR does not support MVLANS or MVR receiver VLANs on a private VLAN (PVLAN).

On non-ELS switches, the MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs.

On ELS switches, the MVR receiver ports are all the interfaces on the MVR receiver VLANs except the multicast router ports; an interface can be configured in both an MVR receiver VLAN and its MVLAN only if it is configured as a multicast router port in both VLANs. ELS EX Series switches support MVR as follows:

- Starting in Junos OS Release 18.3R1, EX4300 switches and Virtual Chassis support MVR. You can configure up to 10 MVLANS on these devices.
- Starting in Junos OS Release 18.4R1, EX2300 and EX3400 switches and Virtual Chassis support MVR. You can configure up to 5 MVLANS on these devices.

- Starting in Junos OS Release 19.4R1, EX4300 multigigabit model (EX4300-48MP) switches and Virtual Chassis support MVR. You can configure up to 10 MVLANs on these devices.



NOTE: MVR has some configuration and operational differences on EX Series switches that use the Enhanced Layer 2 Software (ELS) configuration style compared to MVR on switches that do not support ELS. Where applicable, the following sections explain these differences.

MVR Modes

MVR can operate in two modes: MVR transparent mode and MVR proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network. However, the main difference between the two modes is in how the device sends IGMP reports upstream to the multicast router. The device essentially handles IGMP queries the same way in either mode.

You configure MVR modes differently on non-ELS and ELS switches. Also, on ELS switches, you can associate an MVLAN with some MVR receiver VLANs operating in proxy mode and others operating in transparent mode if you have multicast requirements for both modes in your network.

MVR Transparent Mode

Transparent mode is the default mode when you configure an MVR receiver VLAN, also called a data-forwarding receiver VLAN.



NOTE: On ELS switches, you can explicitly configure transparent mode, although it is also the default setting if you don't configure an MVR receiver mode.

In MVR transparent mode, the device handles IGMP packets destined for both the multicast source VLAN and multicast receiver VLANs similarly to the way that it handles them when MVR is not being used. Without MVR, when a host on a VLAN sends IGMP join and leave messages, the device forwards the messages to all multicast router interfaces in the VLAN. Similarly, when a VLAN receives IGMP queries from its multicast router interfaces, it forwards the queries to all interfaces in the VLAN.

With MVR in transparent mode, the device handles IGMP reports and queries as follows:

- Receives IGMP join and leave messages on MVR receiver VLAN interfaces and forwards them to the multicast router ports on the MVR receiver VLAN.
- Forwards IGMP queries on the MVR receiver VLAN to all MVR receiver ports.
- Forwards IGMP queries received on the MVLAN only to the MVR receiver ports that are in receiver VLANs associated with that MVLAN, even though those ports might not be on the MVLAN itself.



NOTE: Devices in transparent mode only send IGMP reports in the context of the MVR receiver VLAN. In other words, if MVR receiver ports receive an IGMP query from an upstream multicast router on the MVLAN, they only send replies on the MVR receiver VLAN multicast router ports. The upstream router (that sent the queries on the MVLAN) does not receive the replies and does not forward any traffic, so to solve this problem, you must configure static membership. As a result, we recommend that you use MVR proxy mode instead of transparent mode on the device that is closest to the upstream multicast router. See "[MVR Proxy Mode](#)" on page 248.

If a host on a multicast receiver port in the MVR receiver VLAN joins a group, the device adds the appropriate bridging entry on the MVLAN for that group. When the device receives traffic on the MVLAN for that group, it forwards the traffic on that port tagged with the MVLAN tag (even though the port is not in the MVLAN). Likewise, if a host on a multicast receiver port on the MVR receiver VLAN leaves a group, the device deletes the matching bridging entry, and the MVLAN stops forwarding that group's MVR traffic on that port.

When in transparent mode, by default, the device installs bridging entries only on the MVLAN that is the source for the group address, so if the device receives MVR receiver VLAN traffic for that group, the device would not forward the traffic to receiver ports on the MVR receiver VLAN that sent the join message for that group. The device only forwards traffic to MVR receiver interfaces on the MVLAN. To enable MVR receiver VLAN ports to receive traffic forwarded on the MVR receiver VLAN, you can configure the `install` option at the `[edit protocols igmp-snooping vlans vlan-name data-forwarding receiver]` hierarchy level so the device also installs the bridging entries on the MVR receiver VLAN.

MVR Proxy Mode

When you configure MVR in proxy mode, the device acts as an IGMP proxy to the multicast router for MVR group membership requests received on MVR receiver VLANs. That means the device forwards IGMP reports from hosts on MVR receiver VLANs in the context of the MVLAN, and only forwards them to the multicast router ports on the MVLAN. The multicast router receives IGMP reports only on the MVLAN for those MVR receiver hosts.

The device handles IGMP queries in the same way as in transparent mode:

- Forwards IGMP queries received on the MVR receiver VLAN to all MVR receiver ports.
- Forwards IGMP queries received on the MVLAN only to the MVR receiver ports that are in receiver VLANs belonging to that MVLAN, even though those ports might not be on the MVLAN itself.

In proxy mode, for multicast group memberships established in the context of the MVLAN, the device installs bridging entries only on the MVLAN and forwards incoming MVLAN traffic to hosts on the MVR receiver VLANs subscribed to those groups. Proxy mode doesn't support the `install` option that enables the device to also install bridging entries on the MVR receiver VLANs. As a result, when the device

receives traffic on an MVR receiver VLAN, it does not forward the traffic to the hosts on the MVR receiver VLAN because the device does not have bridging entries for those MVR receiver ports on the MVR receiver VLANs.

Proxy Mode on Non-ELS Switches

On non-ELS switches, you configure MVR proxy mode on an MVLAN using the *proxy* statement at the [edit protocols igmp-snooping vlan *vlan-name*] hierarchy level along with other IGMP snooping configuration options.



NOTE: On non-ELS switches, this *proxy* configuration statement only supports MVR proxy mode configuration. General IGMP snooping proxy operation is not supported.

When this option is enabled on non-ELS switches, the device acts as an IGMP proxy for any MVR groups sourced by the MVLAN in both the upstream and downstream directions. In the downstream direction, the device acts as the querier for those multicast groups in the MVR receiver VLANs. In the upstream direction, the device originates the IGMP reports and leave messages, and answers IGMP queries from multicast routers. Configuring this proxy option on an MVLAN automatically enables MVR proxy operation for all MVR receiver VLANs associated with the MVLAN.

Proxy Mode on ELS Switches

On ELS switches, you configure MVR proxy mode on the MVR receiver VLANs. You can configure MVR proxy mode separately from IGMP snooping proxy mode, as follows:

- *IGMP snooping proxy mode*—You can use the *proxy* statement at the [edit protocols igmp-snooping vlan *vlan-name*] hierarchy level on ELS switches to enable IGMP proxy operation with or without MVR configuration. When you configure this option for a VLAN without configuring MVR, the device acts as an IGMP proxy to the multicast router for ports in that VLAN. When you configure this option on an MVLAN, the device acts as an IGMP proxy between the multicast router and hosts in any associated MVR receiver VLANs.



NOTE: You configure this proxy mode on the MVLAN only, not on MVR receiver VLANs.

- *MVR proxy mode*—On ELS switches, you configure MVR proxy mode on an MVR receiver VLAN (rather than on the MVLAN), using the *proxy* option at the [edit igmp-snooping vlan *vlan-name* data-forwarding receiver *mode*] hierarchy level, when you associate the MVR receiver VLAN with an MVLAN. An ELS switch operating in MVR proxy mode for an MVR receiver VLAN acts as an IGMP proxy for that MVR receiver VLAN to the multicast router in the context of the MVLAN.

MVR VLAN Tag Translation

When you configure MVR, the device sends multicast traffic and IGMP queries packets downstream to hosts in the context of the MVLAN by default. The MVLAN tag is included for VLAN-tagged traffic egressing on trunk ports, while traffic egressing on access ports is untagged.

On ELS EX Series switches that support MVR, for VLANs with trunk ports and hosts on a multicast receiver VLAN that expect traffic in the context of that receiver VLAN, you can configure the device to translate the MVLAN tags into the multicast receiver VLAN tags. See the `translate` option at the `[edit protocols igmp-snooping vlans vlan-name data-forwarding receiver]` hierarchy level.

Recommended MVR Configurations in the Access Layer on ELS Switches

Based on the access layer topology of your network, the following sections describe recommended ways you should configure MVR on devices in the access layer to smoothly deliver a single multicast stream to subscribed hosts in multiple VLANs.




NOTE: These sections apply to EX Series switches running Junos OS with the Enhanced Layer 2 Software (ELS) configuration style only.

MVR in a Single-Tier Access Layer Topology

[Figure 30 on page 251](#) shows a device in a single-tier access layer topology. The device is connected to a multicast router in the upstream direction (INTF-1), with host trunk or access ports in the downstream direction connected to multicast receivers in two different VLANs (v10 on INTF-2 and v20 on INTF-3).

Figure 30: MVR in a Single-Tier Access Layer Topology



Without MVR, the upstream interface (INTF-1) acts as a multicast router interface to the upstream router and a trunk port in both VLANs. In this configuration, the upstream router would require two integrated routing and bridging (IRB) interfaces to send two copies of the multicast stream to the device, which then would forward the traffic to the receivers on the two different VLANs on INTF-2 and INTF-3.

With MVR configured as indicated in [Figure 30 on page 251](#), the multicast stream can be sent to receivers in different VLANs in the context of a single MVLAN, and the upstream router only requires one downstream IRB interface on which to send one MVLAN stream to the device.

For MVR to operate smoothly in this topology, we recommend you set up the following elements on the single-tier device as illustrated in [Figure 30 on page 251](#):

- An MVLAN with the device's upstream multicast router interface configured as a trunk port and a multicast router interface in the MVLAN. This upstream interface was already a trunk port and a multicast router port for the receiver VLANs that will be associated with the MVLAN.

[Figure 30 on page 251](#) shows an MVLAN configured on the device, and the upstream interface INTF-1 configured previously as a trunk port and multicast router port in v10 and v20, is subsequently added as a trunk and multicast router port in the MVLAN as well.

- MVR receiver VLANs associated with the MVLAN.

In [Figure 30 on page 251](#), the device is connected to Host 1 on VLAN v10 (using trunk interface INTF-2) and Host 2 on v20 (using access interface INTF-3). VLANs v10 and v20 use INTF-1 as a trunk port and multicast router port in the upstream direction. These VLANs become MVR receiver

VLANs for the MVLAN, with INTF-1 also added as a trunk port and multicast router port in the MVLAN.

- MVR running in proxy mode on the device, so the device processes MVR receiver VLAN IGMP group memberships in the context of the MVLAN. The upstream router sends only one multicast stream on the MVLAN downstream to the device, which is forwarded to hosts on the MVR receiver VLANs that are subscribed to the multicast groups sourced by the MVLAN.

The device in [Figure 30 on page 251](#) is configured in proxy mode and establishes group memberships on the MVLAN for hosts on MVR receiver VLANs v10 and v20. The upstream router in the figure sends only one multicast stream on the MVLAN through INTF-1 to the device, which forwards the traffic to subscribed hosts on MVR receiver VLANs v10 and v20.

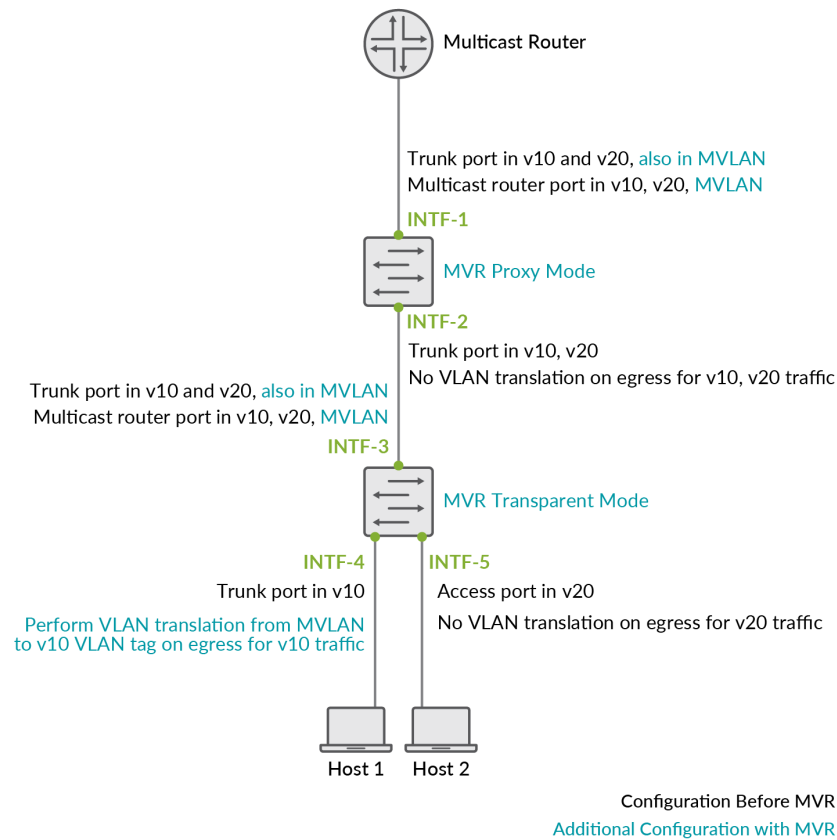
- MVR receiver VLAN tag translation enabled on receiver VLANs that have hosts on trunk ports, so those hosts receive the multicast traffic in the context of their receiver VLANs. Hosts reached by way of access ports receive untagged multicast packets (and don't need MVR VLAN tag translation).

In [Figure 30 on page 251](#), the device has translation enabled on v10 and substitutes the v10 VLAN tag for the mvlan VLAN tag when forwarding the multicast stream on trunk interface INTF-2. The device does not have translation enabled on v20, and forwards untagged multicast packets on access port INTF-3.

MVR in a Multiple-Tier Access Layer Topology

[Figure 31 on page 253](#) shows devices in a two-tier access layer topology. The upper or upstream device is connected to the multicast router in the upstream direction (INTF-1) and to a second device downstream (INTF-2). The lower or downstream device connects to the upstream device (INTF-3), and uses trunk or access ports in the downstream direction to connect to multicast receivers in two different VLANs (v10 on INTF-4 and v20 on INTF-5).

Figure 31: MVR in a Multiple-Tier Access Layer Topology



Without MVR, similar to the single-tier access layer topology, the upper device connects to the upstream multicast router using a multicast router interface that is also a trunk port in both receiver VLANs. The two layers of devices are connected with trunk ports in the receiver VLANs. The lower device has trunk or access ports in the receiver VLANs connected to the multicast receiver hosts. In this configuration, the upstream router must duplicate the multicast stream and use two IRB interfaces to send copies of the same data to the two VLANs. The upstream device also sends duplicate streams downstream for receivers on the two VLANs.

With MVR configured as shown in [Figure 31 on page 253](#), the multicast stream can be sent to receivers in different VLANs in the context of a single MVLAN from the upstream router and through the multiple tiers in the access layer.

For MVR to operate smoothly in this topology, we recommend to set up the following elements on the different tiers of devices in the access layer, as illustrated in [Figure 31 on page 253](#):

- An MVLAN configured on the devices in all tiers in the access layer. The device in the uppermost tier connects to the upstream multicast router with a multicast router interface and a trunk port in the MVLAN. This upstream interface was already a trunk port and a multicast router port for the receiver VLANs that will be associated with the MVLAN.

Figure 31 on page 253 shows an MVLAN configured on all tiers of devices. The upper-tier device is connected to the multicast router using interface INTF-1, configured previously as a trunk port and multicast router port in v10 and v20, and subsequently added to the configuration as a trunk and multicast router port in the MVLAN as well.

- MVR receiver VLANs associated with the MVLAN on the devices in all tiers in the access layer.

In Figure 31 on page 253, the lower-tier device is connected to Host 1 on VLAN v10 (using trunk interface INTF-4) and Host 2 on v20 (using access interface INTF-5). VLANs v10 and v20 use INTF-3 as a trunk port and multicast router port in the upstream direction to the upper-tier device. The upper device connects to the lower device using INTF-2 as a trunk port in the downstream direction to send IGMP queries and forward multicast traffic on v10 and v20. VLANs v10 and v20 are then configured as MVR receiver VLANs for the MVLAN, with INTF-3 also added as a trunk port and multicast router port in the MVLAN. VLANs v10 and v20 are also configured on the upper-tier device as MVR receiver VLANs for the MVLAN.

- MVR running in proxy mode on the device in the uppermost tier for the MVR receiver VLANs, so the device acts as a proxy to the multicast router for group membership requests received on the MVR receiver VLANs. The upstream router sends only one multicast stream on the MVLAN downstream to the device.

In Figure 31 on page 253, the upper-tier device is configured in proxy mode and establishes group memberships on the MVLAN for hosts on MVR receiver VLANs v10 and v20. The upstream router in the figure sends only one multicast stream on the MVLAN, which reaches the upper device through INTF-1. The upper device forwards the stream to the devices in the lower tiers using INTF-2.

- No MVR receiver VLAN tag translation enabled on MVLAN traffic egressing from upper-tier devices. Devices in the intermediate tiers should forward MVLAN traffic downstream in the context of the MVLAN, tagged with the MVLAN tag.

The upper device in the figure does not have translation enabled for either receiver VLAN v10 or v20 for the interface INTF-2 that connects to the lower-tier device.

- MVR running in transparent mode on the devices in the lower tiers of the access layer. The lower devices send IGMP reports upstream in the context of the receiver VLANs because they are operating in transparent mode, and install bridging entries for the MVLAN only, by default, or with the `install` option configured, for both the MVLAN and the MVR receiver VLANs. The uppermost device is running in proxy mode and installs bridging entries for the MVLAN only. The upstream router sends only one multicast stream on the MVLAN downstream toward the receivers, and the traffic is forwarded to the MVR receiver VLANs in the context of the MVLAN, with VLAN tag translation if the `translate` option is enabled (described next).

In Figure 31 on page 253, the lower device is connected to the upper device with INTF-3 as a trunk port and the multicast router port for receiver VLANs v10 and v20. To enable MVR on the lower-tier device, the two MVR receiver VLANs are configured in MVR transparent mode, and INTF-3 is additionally configured to be a trunk port and multicast router port for the MVLAN.

- MVR receiver VLAN tag translation enabled on receiver VLANs on lower-tier devices that have hosts on trunk ports, so those hosts receive the multicast traffic in the context of their receiver VLANs. Hosts reached by way of access ports receive untagged packets, so no VLAN tag translation is needed in that case.

In [Figure 31 on page 253](#), the device has translation enabled on v10 and substitutes the v10 receiver VLAN tag for mvlan's VLAN tag when forwarding the multicast stream on trunk interface INTF-4. The device does not have translation enabled on v20, and forwards untagged multicast packets on access port INTF-5.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.4R1	Starting in Junos OS Release 19.4R1, EX4300 multigigabit model (EX4300-48MP) switches and Virtual Chassis support MVR. You can configure up to 10 MVLANS on these devices.
18.4R1	Starting in Junos OS Release 18.4R1, EX2300 and EX3400 switches and Virtual Chassis support MVR. You can configure up to 5 MVLANS on these devices.
18.3R1	Starting in Junos OS Release 18.3R1, EX4300 switches and Virtual Chassis support MVR. You can configure up to 10 MVLANS on these devices.

RELATED DOCUMENTATION

[Configuring Multicast VLAN Registration on EX Series Switches | 255](#)

[Understanding Multicast VLAN Registration | 245](#)

Understanding FIP Snooping, FBF, and MVR Filter Scalability

Configuring Multicast VLAN Registration on EX Series Switches

IN THIS SECTION

- [Configuring Multicast VLAN Registration on EX Series Switches with ELS | 256](#)
- [Viewing MVLAN and MVR Receiver VLAN Information on EX Series Switches with ELS | 265](#)

● Configuring Multicast VLAN Registration on non-ELS EX Series Switches | 266

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast VLAN (MVLAN) to receive multicast streams from the MVLAN, sharing the MVLAN across multiple VLANs in a Layer 2 network. Hosts remain in their own VLANs for bandwidth and security reasons but are able to receive multicast streams on the MVLAN.

MVR is not enabled by default on switches that support MVR. You must explicitly configure a switch with a *data-forwarding source MVLAN* and associate it with one or more *data-forwarding MVR receiver VLANs*. When you configure one or more VLANs on a switch to be MVR receiver VLANs, you must configure at least one associated source MVLAN. However, you can configure a source MVLAN without associating MVR receiver VLANs with it at the same time.

The overall purpose and benefits of employing MVR are the same on switches that use Enhanced Layer 2 Software (ELS) configuration style and those that do not use ELS. However, there are differences in MVR configuration and operation on the two types of switches.

Configuring Multicast VLAN Registration on EX Series Switches with ELS

The following are configuration frameworks we recommended for MVR to operate smoothly on EX Series switches that support Enhanced Layer 2 Software (ELS) configuration style in single-tier or multiple-tier access layers:

- In an access layer with *a single tier of switches*, where a switch is connected to a multicast router in the upstream direction, and has host trunk or access ports connecting to downstream multicast receivers:
 - Configure MVR on the receiver VLANs to operate in proxy mode.
 - Statically configure the upstream interface to the multicast router as a multicast router port in the MVLAN.
 - Configure the `translate` option on MVR receiver VLANs that have trunk ports, so hosts on those trunk ports receive the multicast packets tagged for their own VLANs.
- In an access layer with *multiple tiers of switches*, with a switch connected upstream to the multicast router and a path through one or more downstream switches to multicast receivers:
 - Configure MVR on the receiver VLANs to operate in proxy mode on the uppermost switch that is directly connected to the upstream multicast router.
 - Configure MVR on the receiver VLANs to operate in transparent mode for the remaining downstream tiers of switches.

- Statically configure a multicast router port to the switch in the upstream direction on each tier for the MVLAN.
- On the lowest tier of MVR switches (connected to receiver hosts), configure MVLAN tag translation for MVR receiver VLANs that have trunk ports, so hosts on those trunk ports receive the multicast stream with the packets tagged with their own VLANs.



NOTE: When enabling MVR on ELS switches, depending on your multicast network requirements, you can have some MVR receiver VLANs configured in proxy mode and some in transparent mode that are associated with the same MVLAN, because the MVR mode setting applies individually to an MVR receiver VLAN. The mode configurations described here are only recommendations for smooth MVR operation in those topologies.

The following constraints apply when configuring MVR on ELS EX Series switches:

- MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.
- You can configure up to 10 MVLANs on an EX4300 or EX4300 multigigabit switch, up to 5 MVLANs on EX2300 and EX3400 switches, and up to a total of 4K MVR receiver VLANs and MVLANs together.
- A VLAN can be configured as either an MVLAN or an MVR receiver VLAN, not both. However, an MVR receiver VLAN can be associated with more than one MVLAN.
- An MVLAN can be the source for only one multicast group subnet, so multiple MVLANs configured on a switch must have unique multicast group subnet ranges.
- You can configure an interface in both an MVR receiver VLAN and its MVLAN only if it is configured as a multicast router port in both VLANs.
- You cannot configure proxy mode with the `install` option to also install forwarding entries on an MVR receiver VLAN. In proxy mode, IGMP reports are sent to the upstream router only in the context of the MVLAN. Multicast sources will not receive IGMP reports on the MVR receiver VLAN, and multicast traffic will not be sent on the MVR receiver VLAN.
- MVR does not support configuring an MVLAN or MVR receiver VLANs on private VLANs (PVLANS).

To configure MVR on ELS EX Series switches that support MVR:

1. Configure a data-forwarding multicast source VLAN as an MVLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan mvlan-name data-forwarding source groups group-subnet
```

For example, configure VLAN `mvlan` as an MVLAN for multicast group subnet `233.252.0.0/8`:

```
[edit protocols igmp-snooping]
user@switch# set vlan mvlan data-forwarding source groups 233.252.0.0/8
```

2. Configure one or more data-forwarding MVR receiver VLANs associated with the source MVLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name data-forwarding receiver source-list mvlan-name
```

For example, configure two MVR receiver VLANs `v10` and `v20` associated with the MVLAN named `mvlan`:

```
[edit protocols igmp-snooping]
user@switch# set vlan v10 data-forwarding receiver source-list mvlan
[edit protocols igmp-snooping]
user@switch# set vlan v20 data-forwarding receiver source-list mvlan
```

3. On a switch in a single-tier topology or on the uppermost switch in a multiple-tier topology (the switch connected to the upstream multicast router), configure each MVR receiver VLAN on the switch to operate in proxy mode:

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name data-forwarding receiver mode proxy
```

For example, configure the two MVR receiver VLANs `v10` and `v20` (associated with the MVLAN named `mvlan`) from the previous step to use proxy mode:

```
[edit protocols igmp-snooping]
user@switch# set vlan v10 data-forwarding receiver mode proxy
[edit protocols igmp-snooping]
user@switch# set vlan v20 data-forwarding receiver mode proxy
```



NOTE: On ELS switches, the MVR mode setting applies to individual MVR receiver VLANs. All MVR receiver VLANs associated with an MVLAN are not required to have the same mode setting. Depending on your multicast network requirements, you might

want to configure some MVR receiver VLANs in proxy mode and others that are associated with the same MVLAN in transparent mode.

4. In a multiple-tier topology, for the remaining switches that are not the uppermost switch, configure each MVR receiver VLAN on each switch to operate in transparent mode. An MVR receiver VLAN operates in transparent mode by default if you do not set the mode explicitly, so this step is optional on these switches.

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name data-forwarding receiver mode transparent
```

For example, configure two MVR receiver VLANs v10 and v20 that are associated with the MVLAN named mvlan to use transparent mode:

```
[edit protocols igmp-snooping]
user@switch# set vlan v10 data-forwarding receiver mode transparent
[edit protocols igmp-snooping]
user@switch# set vlan v20 data-forwarding receiver mode transparent
```

5. Configure a multicast router port in the upstream direction for the MVLAN on the MVR switch in a single-tier topology or on the MVR switch in each tier of a multiple-tier topology:

```
[edit protocols igmp-snooping]
user@switch# set vlan mvlan-name interface interface-name multicast-router-interface
```

For example, configure a multicast router interface ge-0/0/10.0 for the MVLAN named mvlan:

```
[edit protocols igmp-snooping]
user@switch# set vlan mvlan interface ge-0/0/10.0 multicast-router-interface
```

6. On an MVR switch connected to the receiver hosts with trunk or access ports (applies only to the lowest tier in a multiple-tier topology), configure MVLAN tag translation on MVR receiver VLANs that have trunk ports, so hosts on the trunk ports can receive the multicast stream with the packets tagged with their own VLANs:

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name data-forwarding receiver translate
```

For example, a switch connects to receiver hosts on MVR receiver VLAN v10 using a trunk port, but reaches receiver hosts on MVR receiver VLAN v20 on an access port, so configure the MVR translate option only on VLAN v10:

```
[edit protocols igmp-snooping]
user@switch# set vlan v10 data-forwarding receiver translate
```

7. (Optional and applicable only to MVR receiver VLANs configured in transparent mode) Install forwarding entries for an MVR receiver VLAN as well as the MVLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name data-forwarding receiver install
```



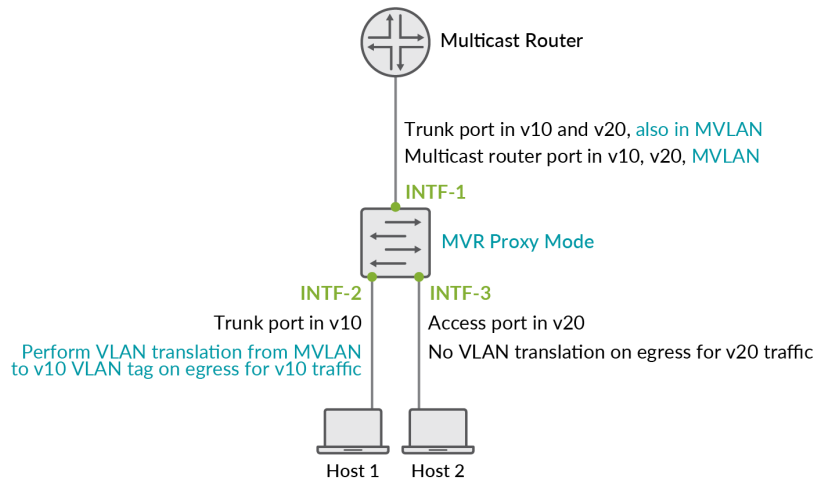
NOTE: This option cannot be configured for an MVR receiver VLAN configured in proxy mode.

For example:

```
[edit protocols igmp-snooping]
user@switch# set vlan v20 data-forwarding receiver install
```

[Figure 32 on page 261](#) illustrates a single-tier access layer topology in which MVR is employed with an MVLAN named `mvlan` and receiver hosts on MVR receiver VLANs `v10` and `v20`. A sample of the recommended MVR configuration for this topology follows the figure.

Figure 32: MVR in a Single-Tier Topology



Configuration Before MVR
Additional Configuration with MVR

8200496

The MVR switch in [Figure 32 on page 261](#) is configured in proxy mode, connects to the upstream multicast router on interface INTF-1, and connects to receiver hosts on v10 using trunk port INTF-2 and on v20 using access port INTF-3. The switch is configured to translate MVLAN tags in the multicast stream into the receiver VLAN tags only for v10 on INTF-2.

```
# Receiver VLAN configuration before configuring MVR
set interfaces INTF-1 unit 0 family ethernet-switching vlan members v10
set interfaces INTF-1 unit 0 family ethernet-switching vlan members v20
set interfaces INTF-1 unit 0 family ethernet-switching interface-mode trunk

set interfaces INTF-2 unit 0 family ethernet-switching vlan members v10
set interfaces INTF-2 unit 0 family ethernet-switching interface-mode trunk
set interfaces INTF-3 unit 0 family ethernet-switching vlan members v20

set vlans v10 vlan-id 10

set vlans v20 vlan-id 20

set protocols igmp-snooping vlan v10
set protocols igmp-snooping vlan v10 interface INTF-1 multicast-router-interface
set protocols igmp-snooping vlan v20
set protocols igmp-snooping vlan v20 interface INTF-1 multicast-router-interface
```

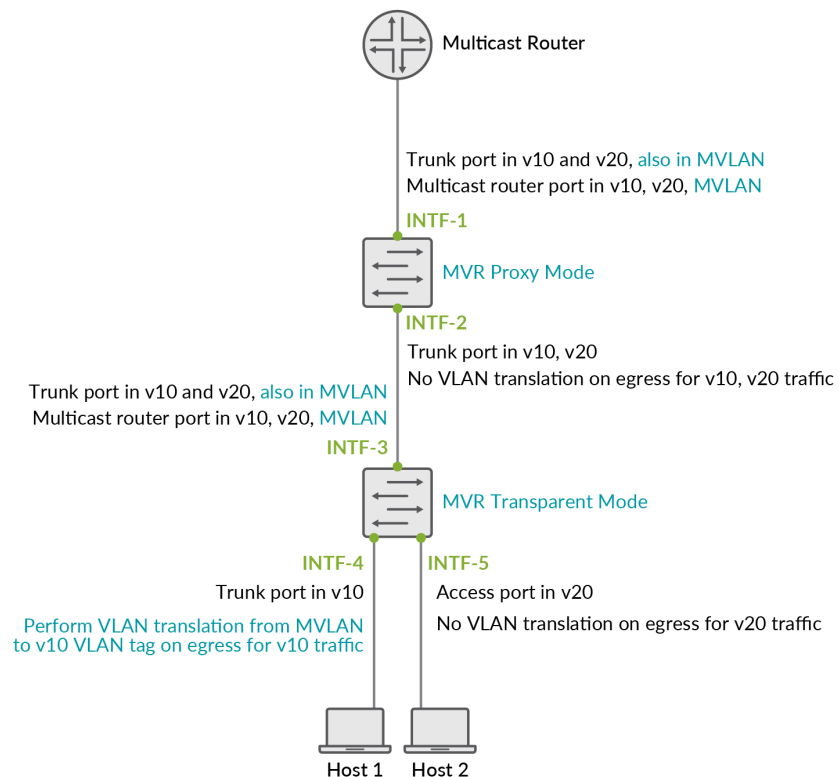
```
# Additional configuration for MVR
set interfaces INTF-1 unit 0 family ethernet-switching vlan members mvlan
set vlans mvlan vlan-id 100
set protocols igmp-snooping vlan mvlan data-forwarding source groups 233.252.0.0/8
set protocols igmp-snooping vlan mvlan interface INTF-1 multicast-router-interface

set protocols igmp-snooping vlan v10 data-forwarding receiver source-list mvlan
set protocols igmp-snooping vlan v10 data-forwarding receiver mode proxy
set protocols igmp-snooping vlan v10 data-forwarding receiver translate

set protocols igmp-snooping vlan v20 data-forwarding receiver source-list mvlan
set protocols igmp-snooping vlan v20 data-forwarding receiver mode proxy
```

Figure 33 on page 262 illustrates a two-tier access layer topology in which MVR is employed with an MVLAN named mvlan, MVR receiver VLANs v10 and v20, and receiver hosts connected to trunk port INTF-4 on v10 and access port INTF-5 on v20. A sample of the recommended MVR configuration for this topology follows the figure.

Figure 33: MVR in a Multiple-Tier Topology



Configuration Before MVR
Additional Configuration with MVR

8200497

The upper switch in [Figure 33 on page 262](#) connects to the upstream multicast router on INTF-1, and the lower switch connects to the upper switch on INTF-3, both configured as trunk ports and multicast router interfaces in the MVLAN. The upper switch is configured in proxy mode and the lower switch is configured in transparent mode for all MVR receiver VLANs. The lower switch is configured to translate MVLAN tags in the multicast stream into the receiver VLAN tags for v10 on INTF-4.

Upper Switch:

```
# Receiver VLAN configuration before configuring MVR
set interfaces INTF-1 unit 0 family ethernet-switching vlan members v10
set interfaces INTF-1 unit 0 family ethernet-switching vlan members v20
set interfaces INTF-1 unit 0 family ethernet-switching interface-mode trunk

set interfaces INTF-2 unit 0 family ethernet-switching vlan members v10
set interfaces INTF-2 unit 0 family ethernet-switching vlan members v20
set interfaces INTF-2 unit 0 family ethernet-switching interface-mode trunk

set vlans v10 vlan-id 10

set vlans v20 vlan-id 20

set protocols igmp-snooping vlan v10
set protocols igmp-snooping vlan v10 interface INTF-1 multicast-router-interface
set protocols igmp-snooping vlan v20
set protocols igmp-snooping vlan v20 interface INTF-1 multicast-router-interface

# Additional configuration for MVR
set interfaces INTF-1 unit 0 family ethernet-switching vlan members mvlan
set vlans mvlan vlan-id 100
set protocols igmp-snooping vlan mvlan data-forwarding source groups 233.252.0.0/8
set protocols igmp-snooping vlan mvlan interface INTF-1 multicast-router-interface

set protocols igmp-snooping vlan v10 data-forwarding receiver source-list mvlan
set protocols igmp-snooping vlan v10 data-forwarding receiver mode proxy

set protocols igmp-snooping vlan v20 data-forwarding receiver source-list m-vlan
set protocols igmp-snooping vlan v20 data-forwarding receiver mode proxy
```

Lower Switch:

```
# Receiver VLAN configuration before configuring MVR
set interfaces INTF-3 unit 0 family ethernet-switching vlan members v10
set interfaces INTF-3 unit 0 family ethernet-switching vlan members v20
set interfaces INTF-3 unit 0 family ethernet-switching interface-mode trunk

set interfaces INTF-4 unit 0 family ethernet-switching vlan members v10
set interfaces INTF-4 unit 0 family ethernet-switching interface-mode trunk

set interfaces INTF-5 unit 0 family ethernet-switching vlan members v20

set vlans v10 vlan-id 10

set vlans v20 vlan-id 20

set protocols igmp-snooping vlan v10
set protocols igmp-snooping vlan v10 interface INTF-3 multicast-router-interface
set protocols igmp-snooping vlan v20
set protocols igmp-snooping vlan v20 interface INTF-3 multicast-router-interface

# Additional configuration for MVR
set interfaces INTF-3 unit 0 family ethernet-switching vlan members mvlan
set protocols igmp-snooping vlan mvlan data-forwarding source groups 233.252.0.0/8
set protocols igmp-snooping vlan mvlan interface INTF-3 multicast-router-interface
set vlans mvlan vlan-id 100

set protocols igmp-snooping vlan v10 data-forwarding receiver source-list mvlan
set protocols igmp-snooping vlan v10 data-forwarding receiver mode transparent
set protocols igmp-snooping vlan v10 data-forwarding receiver translate

set protocols igmp-snooping vlan v20 data-forwarding receiver source-list mvlan
set protocols igmp-snooping vlan v20 data-forwarding receiver mode transparent
```

Viewing MVLAN and MVR Receiver VLAN Information on EX Series Switches with ELS

On EX Series switches with the Enhanced Layer 2 Software (ELS) configuration style that support MVR, you can use the `show igmp snooping data-forwarding` command to view information about the MVLANS and MVR receiver VLANs configured on a switch, as follows:

```
user@host> show igmp snooping data-forwarding
```

```
Instance: default-switch
```

```
Vlan: v2
```

```
Learning-Domain : default
Type             : MVR Source Vlan
Group subnet    : 225.0.0.0/24
Receiver vlans:
  vlan: v1
  vlan: v3
```

```
Vlan: v1
```

```
Learning-Domain : default
Type            : MVR Receiver Vlan
Mode           : PROXY
Egress translate : FALSE
Install route  : FALSE
Source vlans:
  vlan: v2
```

```
Vlan: v3
```

```
Learning-Domain : default
Type            : MVR Receiver Vlan
Mode           : TRANSPARENT
Egress translate : FALSE
Install route  : TRUE
Source vlans:
  vlan: v2
```

MVLANS are listed as Type: MVR Source Vlan with the associated group subnet range and MVR receiver VLANs. MVR receiver VLANs are listed as Type: MVR Receiver Vlan with the associated source MVLANS

and configured options (proxy or transparent mode, VLAN tag translation, and installation of receiver VLAN forwarding entries).

In addition, the `show igmp snooping interface` and `show igmp snooping membership` commands on ELS EX Series switches list MVR receiver VLAN interfaces under both the MVR receiver VLAN and its MVLAN, and display the output field `Data-forwarding receiver: yes` when MVR receiver ports are listed under the MVLAN. This field is not displayed for other interfaces in an MVLAN listed under the MVLAN that are not in MVR receiver VLANs.

Configuring Multicast VLAN Registration on non-ELS EX Series Switches

When you configure MVR on EX Series switches that do not support Enhanced Layer 2 Software (ELS) configuration style, the following constraints apply:

- MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.
- A VLAN can be configured as an MVLAN or an MVR receiver VLAN, but not both. However, an MVR receiver VLAN can be associated with more than one MVLAN.
- An MVLAN can be the source for only one multicast group subnet, so multiple MVLANS configured on a switch must have disjoint multicast group subnets.
- After you configure a VLAN as an MVLAN, that VLAN is no longer available for other uses.
- You cannot enable multicast protocols on VLAN interfaces that are members of MVLANS.
- If you configure an MVLAN in proxy mode, IGMP snooping proxy mode is automatically enabled on all MVR receiver VLANs of this MVLAN. If a VLAN is an MVR receiver VLAN for multiple MVLANS, all of the MVLANS must have proxy mode enabled or all must have proxy mode disabled. You can enable proxy mode only on VLANs that are configured as MVR source VLANs and that are not configured for Q-in-Q tunneling.
- You cannot configure proxy mode with the `install` option to also install forwarding entries for received IGMP packets on an MVR receiver VLAN.

To configure MVR on switches that do not support ELS:

1. Configure the VLAN named `mv0` to be an MVLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 data-forwarding source groups 225.10.0.0/16
```

2. Configure the MVLAN mv0 to be a proxy VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 proxy source-address 10.0.0.1
```

3. Configure the VLAN named v2 to be an MVR receiver VLAN with mv0 as its source:

```
[edit protocols]
user@switch# set igmp-snooping vlan v2 data-forwarding receiver source-vlans mv0
```

4. Install forwarding entries in the MVR receiver VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan v2 data-forwarding receiver install
```

SEE ALSO

| [Understanding Multicast VLAN Registration | 245](#)

RELATED DOCUMENTATION

| [Understanding Multicast VLAN Registration | 245](#)

Example: Configuring Multicast VLAN Registration on EX Series Switches Without ELS

IN THIS SECTION

- [Requirements | 268](#)
- [Overview and Topology | 268](#)
- [Configuration | 271](#)

Multicast VLAN registration (MVR) enables hosts that are not part of a multicast VLAN (MVLAN) to receive multicast streams from the MVLAN, which enable the MVLAN to be shared across the Layer 2 network and eliminate the need to send duplicate multicast streams to each requesting VLAN in the network. Hosts remain in their own VLANs for bandwidth and security reasons.



NOTE: This example describes configuring MVR only on EX Series and QFX Series switches that do not support the Enhanced Layer 2 Software configuration style.

Requirements

This example uses the following hardware and software components:

- One EX Series or QFX Series switch
- Junos OS Release 9.6 or later for EX Series switches or Junos OS Release 12.3 or later for the QFX Series

Before you configure MVR, be sure you have:

- Configured two or more VLANs on the switch. See the task for your platform:
 - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches](#)
 - [Example: Setting Up Bridging with Multiple VLANs on Switches](#) for the QFX Series and EX4600 switch
- Connected the switch to a network that can transmit IPTV multicast streams from a video server.
- Connected a host that is capable of receiving IPTV multicast streams to an interface in one of the VLANs.

Overview and Topology

IN THIS SECTION

- [Topology | 269](#)

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which multicast traffic flows throughout the Layer 2 network. Multicast

traffic can then be selectively forwarded from interfaces on the MVLAN (source ports) to hosts that are connected to interfaces (multicast receiver ports) that are not part of the multicast source VLAN. When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs.

Topology

You can configure MVR to operate in one of two modes: transparent mode (the default mode) or proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

In transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. [Figure 34 on page 270](#) shows how MVR operates in transparent mode.

In proxy mode, the switch acts as a proxy for the IGMP multicast router in the MVLAN for MVR group memberships established in the MVR receiver VLANs and generates and sends IGMP packets into the MVLAN as needed. [Figure 35 on page 271](#) shows how MVR operates in proxy mode.

This example shows how to configure MVR in both transparent mode and proxy mode on an EX Series switch or the QFX Series. The topology includes a video server that is connected to a multicast router, which in turn forwards the IPTV multicast traffic in the MVLAN to the Layer 2 network.

[Figure 34 on page 270](#) shows the MVR topology in transparent mode. Interfaces P1 and P2 on Switch C belong to service VLAN s0 and MVLAN mv0. Interface P4 of Switch C also belongs to service VLAN s0. In the upstream direction of the network, only non-IPTV traffic is being carried in individual customer VLANs of service VLAN s0. VLAN c0 is an example of this type of customer VLAN. IPTV traffic is being carried on MVLAN mv0. If any host on any customer VLAN connected to port P4 requests an MVR stream, Switch C takes the stream from VLAN mv0 and replicates that stream onto port P4 with tag mv0. IPTV traffic, along with other network traffic, flows from port P4 out to the Digital Subscriber Line Access Multiplexer (DSLAM) D1.

Figure 34: MVR Topology in Transparent Mode





Figure 35 on page 271 shows the MVR topology in proxy mode. Interfaces P1 and P2 on Switch C belong to MVLAN mv0 and customer VLAN c0. Interface P4 on Switch C is an access port of customer VLAN c0. In the upstream direction of the network, only non-IPTV traffic is being carried on customer VLAN c0. Any IPTV traffic requested by hosts on VLAN c0 is replicated untagged to port P4 based on streams received in MVLAN mv0. IPTV traffic flows from port P4 out to an IPTV-enabled device in Host H1. Other traffic, such as data and voice traffic, also flows from port P4 to other network devices in Host H1.

Figure 35: MVR Topology in Proxy Mode



For information on VLAN tagging, see the topic for your platform:

- [Understanding Bridging and VLANs on Switches](#)

Configuration

IN THIS SECTION

- Procedure | 272

Procedure

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 protocols igmp-snooping]` hierarchy level.

```
set vlan mv0 data-forwarding source groups 225.10.0.0/16
set vlan v2 data-forwarding receiver source-vlans mv0
set vlan v2 data-forwarding receiver install
set vlan mv0 proxy source-address 10.1.1.1
```

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 MVR:

1. Configure VLAN mv0 to be an MVLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 data-forwarding source groups 225.10.0.0/16
```

2. Configure VLAN v2 to be a multicast receiver VLAN with mv0 as its source:

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver source-vlans mv0
```

3. (Optional) Install forwarding entries in the multicast receiver VLAN v2:

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver install
```

4. (Optional) Configure MVR in proxy mode:

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 proxy source-address 10.1.1.1
```

Results

From configuration mode, confirm your configuration by entering the show command at the [edit protocols igmp-snooping] hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit protocols igmp-snooping]
user@switch# show
vlan mv0 {
  proxy {
    source-address 10.1.1.1;
  }
  data-forwarding {
    source {
      groups 225.10.0.0/16;
    }
  }
}
vlan v2 {
  data-forwarding {
    receiver {
      source-vlans mv0;
      install;
    }
  }
}
```

RELATED DOCUMENTATION

[Configuring Multicast VLAN Registration on EX Series Switches | 255](#)

[Understanding Multicast VLAN Registration | 245](#)

3

PART

Configuring Protocol Independent Multicast

[Understanding PIM | 275](#)

[Configuring PIM Basics | 280](#)

[Routing Content to Densely Clustered Receivers with PIM Dense Mode | 296](#)

[Routing Content to Larger, Sparser Groups with PIM Sparse Mode | 307](#)

[Configuring Designated Routers | 424](#)

[Receiving Content Directly from the Source with SSM | 429](#)

[Minimizing Routing State Information with Bidirectional PIM | 470](#)

[Rapidly Detecting Communication Failures with PIM and the BFD Protocol | 497](#)

[Configuring PIM Options | 515](#)

[Verifying PIM Configurations | 539](#)

Understanding PIM

IN THIS CHAPTER

- [PIM Overview | 275](#)
- [PIM on Aggregated Interfaces | 279](#)

PIM Overview

IN THIS SECTION

- [Basic PIM Network Components | 277](#)

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many) resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

Starting in Junos OS Release 15.2, only PIM version 2 is supported. In the CLI, the command for specifying a version (1 or 2) is removed.

PIMv1 and PIMv2 can coexist on the same routing device and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet

structure. All routing devices connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the routing device interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the routing device processes the PIM message, a routing device can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called “protocol independent” because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.



NOTE: ACX Series routers supports only sparse mode. Dense mode on ACX series is supported only for control multicast groups for auto-discovery of rendezvous point (auto-RP).

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [* ,G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to a downstream routing device unless the downstream routing device has sent an explicit request (by means of a join message) to the rendezvous point (RP) routing device to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.

Starting in Junos OS Release 19.2R1, on SRX300, SRX320, SRX340, SRX345, SRX550, SRX1500, and vSRX Virtual Firewall 2.0 and vSRX Virtual Firewall 3.0 (with 2 vCPUs) Series devices, Protocol Independent Multicast (PIM) using point-to-multipoint (P2MP) mode supports AutoVPN and Auto Discovery VPN in which a new p2mp interface type is introduced for PIM. The p2mp interface tracks all PIM joins per neighbor to ensure multicast forwarding or replication only happens to those neighbors that are in joined state. In addition, the PIM using point-to-multipoint mode supports chassis cluster mode.



NOTE: On all the EX series switches (except EX4300 and EX9200), QFX5100 switches, and OCX series switches, the rate limit is set to 1pps per S,G multicast group to avoid overwhelming the rendezvous point (RP), First hop router (FHR) with PIM-sparse mode (PIM-SM) register messages and cause CPU hogs. This rate limit helps in improving scaling and convergence times by avoiding duplicate packets being trapped, and tunneled to RP in software. (Platform support depends on the Junos OS release in your installation.)

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routing devices build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (*,G) state.
- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routing devices sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a routing device receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the outgoing interface list becomes empty, the routing device sends a PIM prune message upstream.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.



NOTE: On SRX Series Firewalls, PIM does not support upstream and downstream interfaces across different virtual routers in flow mode.

Basic PIM Network Components

PIM dense mode requires only a multicast source and series of multicast-enabled routing devices running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routing devices called *rendezvous points (RPs)* in the network core. These routing devices are where upstream join messages from interested receivers meet downstream traffic from the source of the multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routing devices find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router (BSR)* running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic. PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable. If it is not routable, then the bootstrap router is unable to send bootstrap messages to update the RP domain members. The `show pim bootstrap` command displays only those bootstrap routers that have routable loopback addresses.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any routing device, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable routing device and advertised to the network.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
19.2R1	Starting in Junos OS Release 19.2R1, on SRX300, SRX320, SRX340, SRX345, SRX550, SRX1500, and vSRX Virtual Firewall 2.0 and vSRX Virtual Firewall 3.0 (with 2 vCPUs) Series devices, Protocol Independent Multicast (PIM) using point-to-multipoint (P2MP) mode supports AutoVPN and Auto Discovery VPN in which a new <code>p2mp</code> interface type is introduced for PIM.
15.2	Starting in Junos OS Release 15.2, only PIM version 2 is supported. In the CLI, the command for specifying a version (1 or 2) is removed.

RELATED DOCUMENTATION

| [Supported IP Multicast Protocol Standards](#) | 21

PIM on Aggregated Interfaces

You can configure several Protocol Independent Multicast (PIM) features on an interface regardless of its PIM mode (bidirectional, sparse, dense, or sparse-dense mode).



NOTE: ACX Series routers supports only sparse mode. Dense mode on ACX series is supported only for control multicast groups for auto-discovery of rendezvous point (auto-RP).

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream of replicated packets in a load-sharing fashion. The multicast aggregate interface is “expanded” into its constituent interfaces in the next-hop database.

RELATED DOCUMENTATION

| [Junos OS Network Interfaces Library for Routing Devices](#)

Configuring PIM Basics

IN THIS CHAPTER

- [Configuring Different PIM Modes | 280](#)
- [Configuring Multiple Instances of PIM | 282](#)
- [Changing the PIM Version | 282](#)
- [Optimizing the Number of Multicast Flows on QFabric Systems | 283](#)
- [Modifying the PIM Hello Interval | 283](#)
- [Preserving Multicast Performance by Disabling Response to the ping Utility | 285](#)
- [Configuring PIM Trace Options | 286](#)
- [Configuring BFD for PIM | 289](#)
- [Configuring BFD Authentication for PIM | 291](#)

Configuring Different PIM Modes

SUMMARY

Configuring different PIM modes in Junos OS.

PIM operates in four modes. The PIM mode you choose determines the PIM configuration properties, therefore you first must decide which mode PIM operates in your network. Each mode has distinct operating advantages in different network environments.

Sparse Mode

In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to a downstream routing device unless the downstream routing device has sent an explicit request (by means of a join message) to the rendezvous point (RP) routing

device to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common. Refer to ["Examples: Configuring PIM Sparse Mode" on page 311](#) to learn how to configure PIM in sparse mode.

Bidirectional Mode

Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routing devices build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (*,G) state.

Refer to ["Example: Configuring Bidirectional PIM" on page 470](#) to learn how to configure PIM in bidirectional mode.

Dense Mode

Unlike sparse mode and bidirectional mode, in which data is forwarded only to routing devices sending an explicit PIM join request, dense mode implements a flood-and-prune mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a routing device receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the outgoing interface list becomes empty, the routing device sends a PIM prune message upstream. Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

Refer to ["Configuring PIM Dense Mode" on page 299](#) to learn how to configure PIM in dense mode.

Sparse-Dense Mode

Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as "dense" is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as "sparse" is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

Refer to ["Configuring PIM Sparse-Dense Mode" on page 304](#) to learn how to configure PIM in sparse-dense mode.

RELATED DOCUMENTATION

No Link Title

Configuring Multiple Instances of PIM

PIM instances are supported only for VRF instance types. You can configure multiple instances of PIM to support multicast over VPNs.

To configure multiple instances of PIM, include the following statements:

```
routing-instances {
  routing-instance-name {
    interface interface-name;
    instance-type vrf;
    protocols {
      pim {
        ... pim-configuration ...
      }
    }
  }
}
```

You can include the statements at the following hierarchy levels:

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

RELATED DOCUMENTATION

[Junos OS Multicast Protocols User Guide](#)

[Junos OS VPNs Library for Routing Devices](#)

Changing the PIM Version

Starting in Junos OS Release 15.2, it is no longer necessary to configure the PIM version. Support for PIM version 1 has been removed and the remaining, default, version is PIM 2.

PIM version 2 is the default for both rendezvous point (RP) mode (at the [edit protocols pim rp static address *address*] hierarchy level) and for interface mode (at the [edit protocols pim interface *interface-name*] hierarchy level).

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.2	Starting in Junos OS Release 15.2, it is no longer necessary to configure the PIM version.

Optimizing the Number of Multicast Flows on QFabric Systems

Because of the distributed nature of QFabric systems, the default configuration does not allow the maximum number of supported Layer 3 multicast flows to be created. To allow a QFabric system to create the maximum number of supported flows, configure the following statement:

```
set fabric routing-options multicast fabric-optimized-distribution
```

After configuring this statement, you must reboot the QFabric Director group to make the change take effect.

Modifying the PIM Hello Interval

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]
user@host# set hello-interval 255
```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the `show pim neighbors detail` command.

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-3/0/2.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

RELATED DOCUMENTATION

| *show pim neighbors*

Preserving Multicast Performance by Disabling Response to the ping Utility

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routers because the stream of packets is replicated multiple times.

You can disable the router's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the router's response to ping packets sent to multicast addresses:

1. Include the `no-multicast-echo` statement:

```
[edit system]
user@host# set no-multicast-echo
```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the `show system statistics icmp` command.

```
user@host> show system statistics icmp
```

```
icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
100 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
echo: 21
21 message responses generated
```

RELATED DOCUMENTATION

Disable the Routing Engine Response to Multicast Ping Packets

show system statistics icmp

Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Table 11: Flags to configure PIM trace options

Flag	Description
all	Trace all operations.
assert	Trace assert messages, which are used to resolve which of the parallel routers connected to a multiaccess LAN is responsible for forwarding packets to the LAN.
autorp	Trace bootstrap, RP, and auto-RP messages.
bidirectional-df-election	Trace bidirectional PIM designated-forwarder (DF) election events.
bootstrap	Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap router and are forwarded, hop by hop, to all routers in that domain.
general	Trace general events.
graft	Trace graft and graft acknowledgment messages.
hello	Trace hello packets, which are sent so that neighboring routers can discover one another.

Table 11: Flags to configure PIM trace options *(Continued)*

Flag	Description
join	Trace join messages, which are sent to join a branch onto the multicast distribution tree.
mdt	Trace messages related to multicast data tunnels.
normal	Trace normal events.
nsr-synchronization	Trace nonstop routing synchronization events
packets	Trace all PIM packets.
policy	Trace poison-route-reverse packets.
prune	Trace prune messages, which are sent to prune a branch off the multicast distribution tree.
register	Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.
route	Trace routing information.
rp	Trace candidate RP advertisements.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [routing-options hierarchy level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the PIM trace file.

```
[edit protocols pim traceoptions]
user@host# set file pim-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols pim traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols pim traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols pim traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match "Rx V1 Require V2"
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

RELATED DOCUMENTATION

[PIM Overview | 275](#)

[Tracing and Logging Junos OS Operations](#)

Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```
[edit protocols pim]
user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the `show bfd session` command.

RELATED DOCUMENTATION

| `show bfd session`

Configuring BFD Authentication for PIM

IN THIS SECTION

- [Configuring BFD Authentication Parameters | 291](#)
- [Viewing Authentication Information for BFD Sessions | 293](#)

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

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication algorithm
keyed-sha-1
```



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

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

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

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication keychain
bfd-pim
```



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

3. Specify the unique security authentication information for BFD sessions:
 - The matching keychain name as specified in Step 2.
 - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
 - The secret data used to allow access to the session.

- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret $ABC123$/ start-time
2009-06-14.10:00:00
```



NOTE: Security Authentication Keychain is not supported on SRX Series Firewalls.

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

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication loose-
check
```

5. (Optional) View your configuration by using the `show bfd session detail` or `show bfd session extensive` command.
6. Repeat these steps to configure the other end of the BFD session.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the `show bfd session detail` and `show bfd session extensive` commands.

The following example shows BFD authentication configured for the **ge-0/1/5** interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key **1** contains the secret data “**\$ABC123/**” and a start time of June 1, 2009, at 9:46:02 AM PST. Key **2** contains the secret data “**\$ABC123/**” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface ge-0/1/5 {
  family inet {
    bfd-liveness-detection {
      authentication {
        key-chain bfd-pim;
        algorithm keyed-sha-1;
      }
    }
  }
}
```

```

}
[edit security]
authentication key-chains {
  key-chain bfd-pim {
    key 1 {
      secret "$ABC123/";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$ABC123/";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
}

```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the `show bfd session detail` command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the `show bfd session extensive` command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.0.2.2	Up	ge-0/1/5.0	0.900	0.300	3

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**
 Session up time 3d 00:34
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated

show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.0.2.2	Up	ge-0/1/5.0	0.900	0.300	3

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**


```

keychain bfd-pim, algo keyed-sha-1, mode strict
Session up time 00:04:42
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.300, min slow interval 1.000
Adaptive async TX interval 0.300, RX interval 0.300
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
Local discriminator 2, remote discriminator 2
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict

```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
9.6	Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported.

RELATED DOCUMENTATION

[Understanding Bidirectional Forwarding Detection Authentication for PIM | 497](#)

[Configuring BFD for PIM | 500](#)

authentication-key-chains

bfd-liveness-detection

show bfd session

Routing Content to Densely Clustered Receivers with PIM Dense Mode

IN THIS CHAPTER

- [Understanding PIM Dense Mode | 296](#)
- [Understanding PIM Sparse-Dense Mode | 298](#)
- [Mixing PIM Sparse and Dense Modes | 299](#)
- [Configuring PIM Dense Mode | 299](#)
- [Configuring PIM Sparse-Dense Mode | 304](#)

Understanding PIM Dense Mode

PIM dense mode is less sophisticated than PIM sparse mode. PIM dense mode is useful for multicast LAN applications, the main environment for all dense mode protocols.

PIM dense mode implements the same flood-and-prune mechanism that DVMRP and other dense mode routing protocols employ. The main difference between DVMRP and PIM dense mode is that PIM dense mode introduces the concept of protocol independence. PIM dense mode can use the routing table populated by any underlying unicast routing protocol to perform reverse-path-forwarding (RPF) checks.

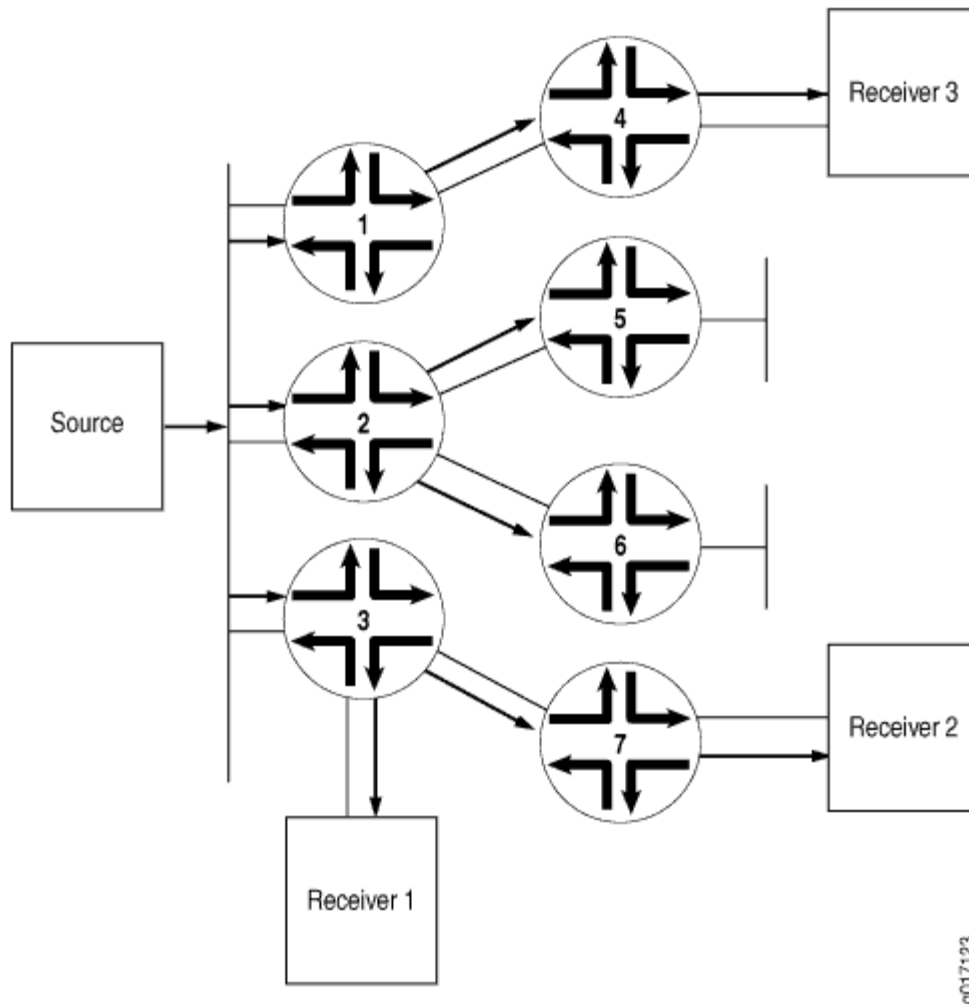
Internet service providers (ISPs) typically appreciate the ability to use any underlying unicast routing protocol with PIM dense mode because they do not need to introduce and manage a separate routing protocol just for RPF checks. While unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in IS-IS (M-IS-IS) were later employed to build special tables to perform RPF checks, PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by OSPF, IS-IS, BGP, and so on, or PIM dense mode can be configured to use a special multicast RPF table populated by MBGP or M-IS-IS when performing RPF checks.

Unlike sparse mode, in which data is forwarded only to routing devices sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is

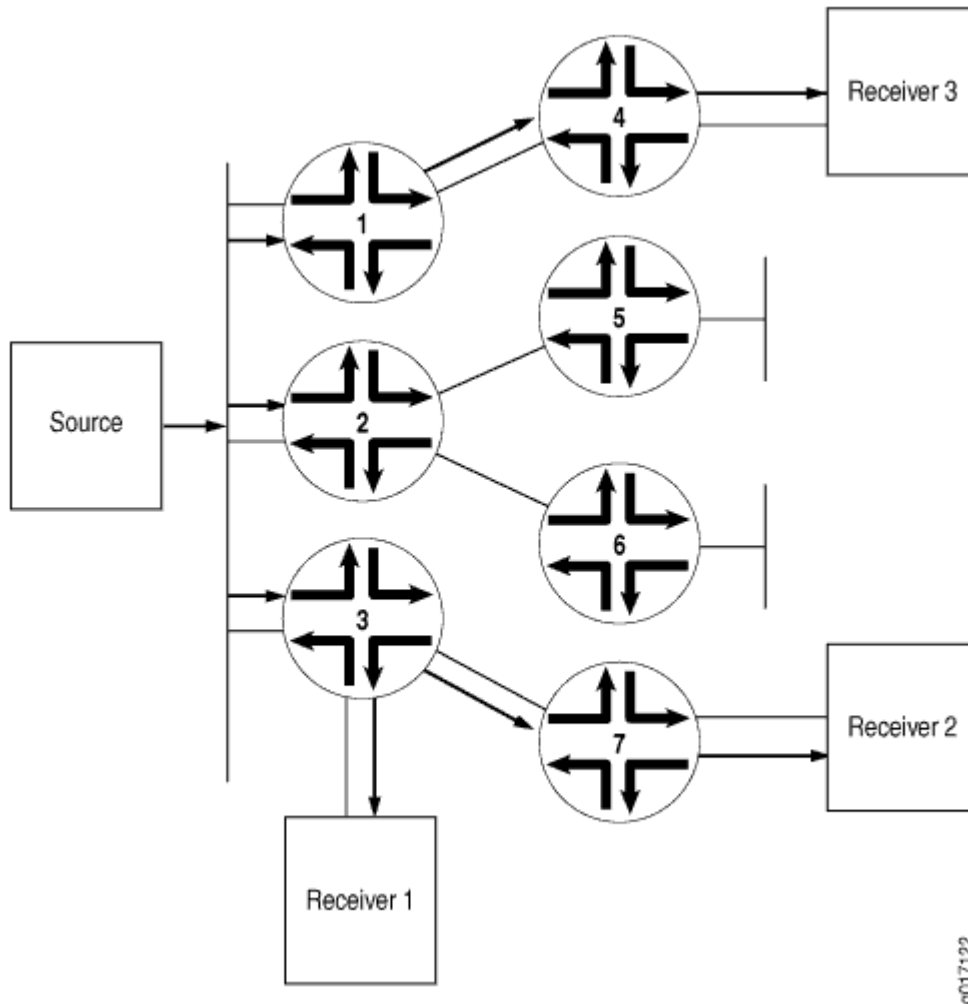
no RP. A routing device receives the multicast data on the interface closest to the source, then forwards the traffic to all other interfaces (see [Figure 36 on page 297](#)).

Figure 36: Multicast Traffic Flooded from the Source Using PIM Dense Mode



Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the OIL becomes empty, the routing device sends a prune message upstream to stop delivery of multicast traffic (see [Figure 37 on page 298](#)).

Figure 37: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic



Understanding PIM Sparse-Dense Mode

Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense-mode rules. A group specified as sparse is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules.

For information about PIM sparse-mode and PIM dense-mode rules, see "[Understanding PIM Sparse Mode](#)" on page 307 and "[Understanding PIM Dense Mode](#)" on page 296.

RELATED DOCUMENTATION

[Understanding PIM Sparse Mode | 307](#)

[Understanding PIM Dense Mode | 296](#)

Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM source-specific multicast (SSM) on the same network, the same routing device, and even the same interface. This is because modes are effectively tied to multicast groups, an IP multicast group address must be unique for a particular group's traffic, and scoping limits enforce the division between potential or actual overlaps.



NOTE: PIM sparse mode was capable of forming shortest-path trees (SPTs) already. Changes to PIM sparse mode to support PIM SSM mainly involved defining behavior in the SSM address range, because shared-tree behavior is prohibited for groups in the SSM address range.

A multicast routing device employing sparse-dense mode is a good example of mixing PIM modes on the same network or routing device or interface. Dense modes are easy to support because of the flooding, but scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

Configuring PIM Dense Mode

IN THIS SECTION

- [Understanding PIM Dense Mode | 299](#)
- [Configuring PIM Dense Mode Properties | 302](#)

Understanding PIM Dense Mode

PIM dense mode is less sophisticated than PIM sparse mode. PIM dense mode is useful for multicast LAN applications, the main environment for all dense mode protocols.

PIM dense mode implements the same flood-and-prune mechanism that DVMRP and other dense mode routing protocols employ. The main difference between DVMRP and PIM dense mode is that PIM dense

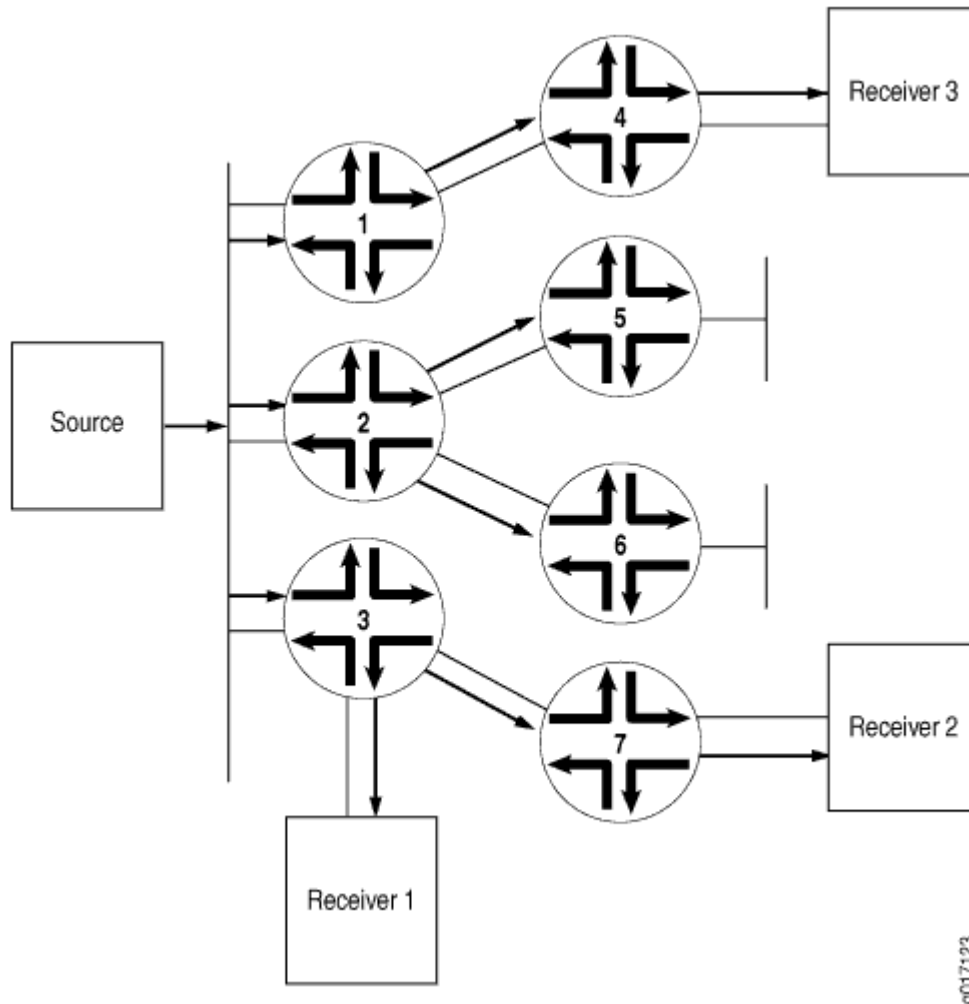
mode introduces the concept of protocol independence. PIM dense mode can use the routing table populated by any underlying unicast routing protocol to perform reverse-path-forwarding (RPF) checks.

Internet service providers (ISPs) typically appreciate the ability to use any underlying unicast routing protocol with PIM dense mode because they do not need to introduce and manage a separate routing protocol just for RPF checks. While unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in IS-IS (M-IS-IS) were later employed to build special tables to perform RPF checks, PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by OSPF, IS-IS, BGP, and so on, or PIM dense mode can be configured to use a special multicast RPF table populated by MBGP or M-IS-IS when performing RPF checks.

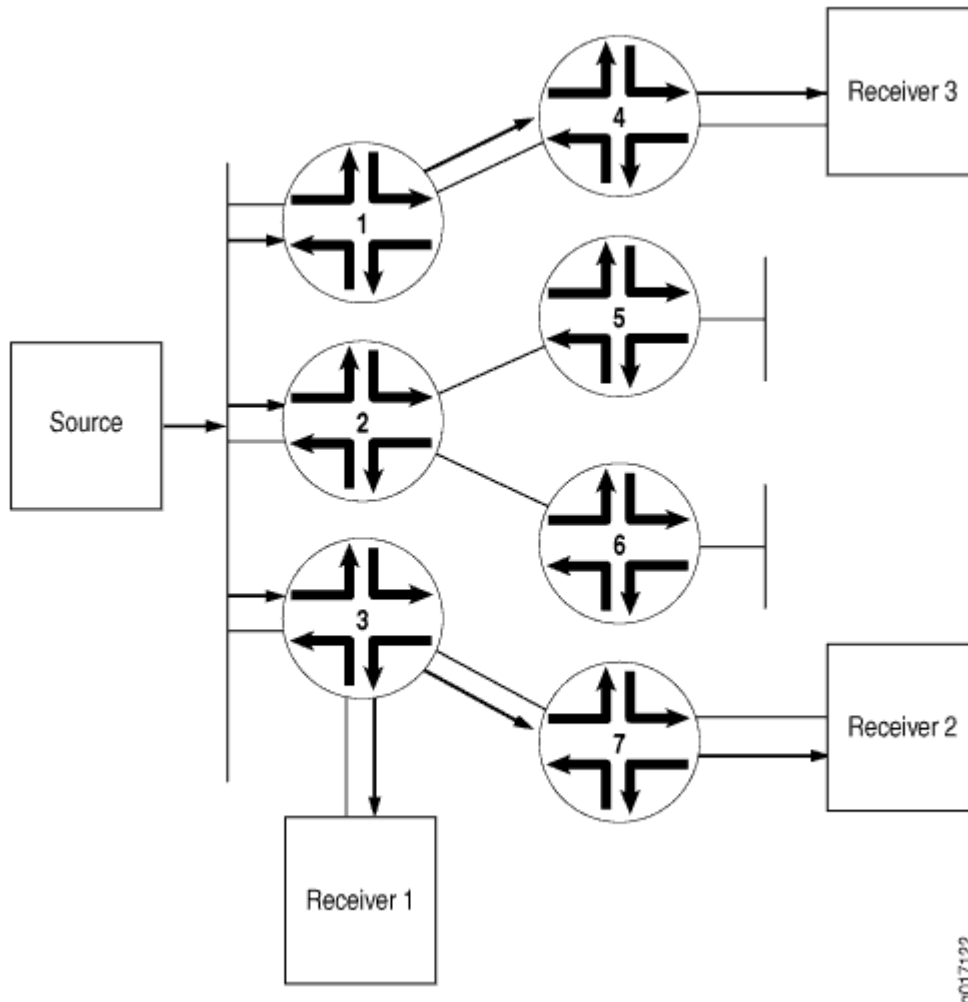
Unlike sparse mode, in which data is forwarded only to routing devices sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is no RP. A routing device receives the multicast data on the interface closest to the source, then forwards the traffic to all other interfaces (see [Figure 38 on page 301](#)).

Figure 38: Multicast Traffic Flooded from the Source Using PIM Dense Mode



Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the OIL becomes empty, the routing device sends a prune message upstream to stop delivery of multicast traffic (see [Figure 39 on page 302](#)).

Figure 39: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic



Configuring PIM Dense Mode Properties

In PIM dense mode (PIM-DM), the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is flooded with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). LANs are appropriate networks for dense-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM dense mode globally or for a routing instance. This example shows how to configure the routing instance and how to specify that PIM dense mode use **inet.2** as its RPF routing table instead of **inet.0**.

To configure the router properties for PIM dense mode:

1. (Optional) Create an IPv4 routing table group so that interface routes are installed into two routing tables, **inet.0** and **inet.2**.

```
[edit routing-options rib-groups]
user@host# set pim-rg export-rib inet.0
user@host# set pim-rg import-rib [ inet.0 inet.2 ]
```

2. (Optional) Associate the routing table group with a PIM routing instance.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set rib-group inet pim-rg
```

3. Configure the PIM interface. If you do not specify any interfaces, PIM is enabled on all router interfaces. Generally, you specify interface names only if you are disabling PIM on certain interfaces.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set interface (Protocols PIM) fe-0/0/1.0 mode dense
```



NOTE: You cannot configure both PIM and Distance Vector Multicast Routing Protocol (DVMRP) in forwarding mode on the same interface. You can configure PIM on the same interface only if you configured DVMRP in unicast-routing mode.

4. Monitor the operation of PIM dense mode by running the `show pim interfaces`, `show pim join`, `show pim neighbors`, and `show pim statistics` commands.

SEE ALSO

[Understanding PIM Dense Mode | 296](#)

[Example: Configuring a Dedicated PIM RPF Routing Table | 1152](#)

RELATED DOCUMENTATION

[Configuring PIM Sparse-Dense Mode | 304](#)

[Configuring Basic PIM Settings](#)

Configuring PIM Sparse-Dense Mode

IN THIS SECTION

- [Understanding PIM Sparse-Dense Mode | 304](#)
- [Mixing PIM Sparse and Dense Modes | 304](#)
- [Configuring PIM Sparse-Dense Mode Properties | 305](#)

Understanding PIM Sparse-Dense Mode

Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense-mode rules. A group specified as sparse is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules.

For information about PIM sparse-mode and PIM dense-mode rules, see "[Understanding PIM Sparse Mode](#)" on page 307 and "[Understanding PIM Dense Mode](#)" on page 296.

SEE ALSO

[Understanding PIM Sparse Mode | 307](#)

[Understanding PIM Dense Mode | 296](#)

Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM source-specific multicast (SSM) on the same network, the same routing device, and even the same interface. This is because modes are effectively tied to multicast groups, an IP multicast group address must be unique for a particular group's traffic, and scoping limits enforce the division between potential or actual overlaps.



NOTE: PIM sparse mode was capable of forming shortest-path trees (SPTs) already. Changes to PIM sparse mode to support PIM SSM mainly involved defining behavior in the SSM address range, because shared-tree behavior is prohibited for groups in the SSM address range.

A multicast routing device employing sparse-dense mode is a good example of mixing PIM modes on the same network or routing device or interface. Dense modes are easy to support because of the flooding, but scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

Configuring PIM Sparse-Dense Mode Properties

Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM sparse-dense mode globally or for a routing instance. This example shows how to configure PIM sparse-dense mode globally on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode.

To configure the router properties for PIM sparse-dense mode:

1. Configure the dense-mode groups.

```
[protocols pim]
user@host# set dense-groups 224.0.1.39
user@host# set dense-groups 224.0.1.40
```

2. Configure all interfaces on the routing device to use sparse-dense mode. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the disable statement for that interface.

```
[edit protocols pim]
user@host# set interface (Protocols PIM) all mode sparse-dense
user@host# set interface (Protocols PIM) fxp0.0 disable
```

3. Monitor the operation of PIM sparse-dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

SEE ALSO

| [Understanding PIM Sparse-Dense Mode](#) | 298

RELATED DOCUMENTATION

[Configuring PIM Dense Mode | 299](#)

[Configuring Basic PIM Settings](#)

Routing Content to Larger, Sparser Groups with PIM Sparse Mode

IN THIS CHAPTER

- Understanding PIM Sparse Mode | 307
- Examples: Configuring PIM Sparse Mode | 311
- Configuring Static RP | 342
- Example: Configuring Anycast RP | 352
- Configuring PIM Bootstrap Router | 365
- Understanding PIM Auto-RP | 371
- Configuring All PIM Anycast Non-RP Routers | 371
- Configuring a PIM Anycast RP Router with MSDP | 372
- Configuring Embedded RP | 373
- Configuring PIM Filtering | 376
- Examples: Configuring PIM RPT and SPT Cutover | 397
- Disabling PIM | 419

Understanding PIM Sparse Mode

IN THIS SECTION

- Rendezvous Point | 309
- RP Mapping Options | 310

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit

join request, an RPF check is triggered. A (*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each router's RPF interface until it reaches the RP. The RP router receives the (*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.



NOTE: State—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and * represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).



NOTE: If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.
- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (*,G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).


PIM sparse mode has standard features for all of these issues.

Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in [Figure 40 on page 310](#), the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

Figure 40: Rendezvous Point As Part of the RPT and SPT



The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

RELATED DOCUMENTATION

[Understanding Static RP | 342](#)

[Understanding RP Mapping with Anycast RP | 353](#)

[Understanding the PIM Bootstrap Router | 365](#)

[Understanding PIM Auto-RP | 371](#)

Examples: Configuring PIM Sparse Mode

IN THIS SECTION

- [Understanding PIM Sparse Mode | 311](#)
- [Understanding Designated Routers | 314](#)
- [Tunnel Services PICs and Multicast | 315](#)
- [Enabling PIM Sparse Mode | 316](#)
- [Configuring PIM Join Load Balancing | 317](#)
- [Modifying the Join State Timeout | 321](#)
- [Example: Enabling Join Suppression | 322](#)
- [Example: Configuring PIM Sparse Mode over an IPsec VPN | 328](#)
- [Example: Configuring Multicast for Virtual Routers with IPv6 Interfaces | 335](#)

Understanding PIM Sparse Mode

IN THIS SECTION

- [Rendezvous Point | 313](#)
- [RP Mapping Options | 314](#)

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A (*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each router's RPF interface until it reaches the RP. The RP router receives the (*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.



NOTE: State—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and * represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).



NOTE: If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.

- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (*,G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).


PIM sparse mode has standard features for all of these issues.

Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in [Figure 41 on page 313](#), the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

Figure 41: Rendezvous Point As Part of the RPT and SPT



The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

SEE ALSO

[Understanding Static RP | 342](#)

[Understanding RP Mapping with Anycast RP | 353](#)

[Understanding the PIM Bootstrap Router | 365](#)

[Understanding PIM Auto-RP | 371](#)

Understanding Designated Routers

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers (DRs) to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.



NOTE: DR priority is specific to PIM sparse mode; as per RFC 3973, PIM DR priority cannot be configured explicitly in PIM Dense Mode (PIM-DM) in IGMPv2 – PIM-DM only support DRs with IGMPv1.

Tunnel Services PICs and Multicast

On Juniper Networks routers, data packets are encapsulated and de-encapsulated into tunnels by means of hardware and not the software running on the router processor. The hardware used to create tunnel interfaces on M Series and T Series routers is a Tunnel Services PIC. If Juniper Networks M Series Multiservice Edge Routers and Juniper Networks T Series Core Routers are configured as rendezvous points or IP version 4 (IPv4) PIM sparse-mode DRs connected to a source, a Tunnel Services PIC is required. Juniper Networks MX Series Ethernet Services Routers do not require Tunnel Services PICs. However, on MX Series routers, you must enable tunnel services with the `tunnel-services` statement on one or more online FPC and PIC combinations at the `[edit chassis fpc number pic number]` hierarchy level.



CAUTION: For redundancy, we strongly recommend that each routing device has multiple Tunnel Services PICs. In the case of MX Series routers, the recommendation is to configure multiple `tunnel-services` statements.

We also recommend that the Tunnel PICs be installed (or configured) on different FPCs. If you have only one Tunnel PIC or if you have multiple Tunnel PICs installed on a single FPC and then that FPC is removed, the multicast session will not come up. Having redundant Tunnel PICs on separate FPCs can help ensure that at least one Tunnel PIC is available and that multicast will continue working.

On MX Series routers, the redundant configuration looks like the following example:

```
[edit chassis]
user@mx-host# set fpc 1 pic 0 tunnel-services bandwidth 1g
user@mx-host# set fpc 2 pic 0 tunnel-services bandwidth 1g
```

In PIM sparse mode, the source DR takes the initial multicast packets and encapsulates them in PIM register messages. The source DR then unicasts the packets to the PIM sparse-mode RP router, where the PIM register message is de-encapsulated.

When a router is configured as a PIM sparse-mode RP router (by specifying an address using the `address` statement at the `[edit protocols pim rp local]` hierarchy level) and a Tunnel PIC is present on the router, a PIM register de-encapsulation interface, or **pd** interface, is automatically created. The **pd** interface receives PIM register messages and de-encapsulates them by means of the hardware.

If PIM sparse mode is enabled and a Tunnel Services PIC is present on the router, a PIM register encapsulation interface (**pe** interface) is automatically created for each RP address. The **pe** interface is

used to encapsulate source data packets and send the packets to RP addresses on the PIM DR and the PIM RP. The **pe** interface receives PIM register messages and encapsulates the packets by means of the hardware.

Do not confuse the configurable **pe** and **pd** hardware interfaces with the nonconfigurable **pime** and **pimd** software interfaces. Both pairs encapsulate and de-encapsulate multicast packets, and are created automatically. However, the **pe** and **pd** interfaces appear only if a Tunnel Services PIC is present. The **pime** and **pimd** interfaces are not useful in situations requiring the **pe** and **pd** interfaces.

If the source DR is the RP, then there is no need for PIM register messages and consequently no need for a Tunnel Services PIC.

When PIM sparse mode is used with IP version 6 (IPv6), a Tunnel PIC is required on the RP, but not on the IPv6 PIM DR. The lack of a Tunnel PIC requirement on the IPv6 DR applies only to IPv6 PIM sparse mode and is not to be confused with IPv4 PIM sparse-mode requirements.

[Table 12 on page 316](#) shows the complete matrix of IPv4 and IPv6 PIM Tunnel PIC requirements.

Table 12: Tunnel PIC Requirements for IPv4 and IPv6 Multicast

IP Version	Tunnel PIC on RP	Tunnel PIC on DR
IPv4	Yes	Yes
IPv6	Yes	No

Enabling PIM Sparse Mode

In PIM sparse mode (PIM-SM), the assumption is that very few of the possible receivers want packets from a source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) a desire for the traffic. WANs are appropriate networks for sparse-mode operation.

Starting in Junos OS Release 16.1, PIM is disabled by default. When you enable PIM, it operates in sparse mode by default. You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After you enable PIM, by default, IGMP version 2 is also enabled.

Junos OS uses PIM version 2 for both rendezvous point (RP) mode (at the [edit protocols pim rp static address *address*] hierarchy level) and interface mode (at the [edit protocols pim interface *interface-name*] hierarchy level).

All systems on a subnet must run the same version of PIM.

You can configure PIM sparse mode globally or for a routing instance. This example shows how to configure PIM sparse mode globally on all interfaces. It also shows how to configure a static RP router and how to configure the non-RP routers.

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

```
[edit protocols pim]
user@host# set rp local family inet address 192.168.3.253
```

2. Configure the RP router interfaces. When configuring all interfaces, exclude the **fxp0.0** management interface by including the `disable` statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse
user@host# set interface fxp0.0 disable
```

3. Configure the non-RP routers. Include the following configuration on all of the non-RP routers.

```
[edit protocols pim]
user@host# set rp static address 192.168.3.253
user@host# set interface all mode sparse
user@host# set interface fxp0.0 disable
```

4. Monitor the operation of PIM sparse mode.

- **show pim interfaces**
- **show pim join**
- **show pim neighbors**
- **show pim rps**

SEE ALSO

| [Understanding PIM Sparse Mode](#) | 307

Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the

join message. This interface is also used for all downstream traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs) and multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, the Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the `show pim join extensive` or `show pim source` commands.

```
user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
```



```

Source: *
RP: 10.255.245.6
Flags: sparse,rptree,wildcard
Upstream interface: t1-0/2/3.0
Upstream neighbor: 192.168.38.57
Upstream state: Join to RP
Downstream neighbors:
  Interface: t1-0/2/1.0
    192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
Source: *
RP: 10.255.245.6
Flags: sparse,rptree,wildcard
Upstream interface: so-0/3/0.0
Upstream neighbor: 192.168.38.47
Upstream state: Join to RP
Downstream neighbors:
  Interface: t1-0/2/3.0
    192.168.38.16 State: JOIN Flags; SRW Timeout: 164

```

Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM sparse mode and join load balancing.

```

[edit protocols pim ]
user@host# set interface all mode sparse version 2
user@host# set join-load-balance

```

3. Then configure the static address of the RP.

```

[edit protocols pim rp]
user@host# set static address 10.10.10.1

```

4. Monitor the operation.

If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the `show pim interfaces` command.

```

user@host> show pim interfaces
Instance: PIM.master

Name          Stat Mode      IP V State NbrCnt JoinCnt  DR address
lo0.0         Up   Sparse       4 2 DR      0      0    10.255.168.58
pe-1/2/0.32769 Up   Sparse       4 2 P2P     0      0
so-0/3/0.0    Up   Sparse       4 2 P2P     1      1
t1-0/2/1.0    Up   Sparse       4 2 P2P     1      0
t1-0/2/3.0    Up   Sparse       4 2 P2P     1      1
lo0.0         Up   Sparse       6 2 DR      0      0    fe80::2a0:a5ff:4b7

```

Note that the two equal-cost paths shown by the `show pim interfaces` command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the `show pim neighbors` detail output:

```

user@host> show pim neighbors detail
Interface: so-0/3/0.0

Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1689116164
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
  BFD: Disabled
  Hello Option Holdtime: 105 seconds 102 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 792890329
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: t1-0/2/3.0

Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 678582286

```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1854475503
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the interface or neighbor is restored, prior joins are not redistributed. The `clear pim join-distribution` command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

SEE ALSO

Load Balancing in Layer 3 VPNs

show pim interfaces

show pim neighbors

show pim source

clear pim join-distribution

Modifying the Join State Timeout

This section describes how to configure the join state timeout.

A downstream router periodically sends join messages to refresh the join state on the upstream router. If the join state is not refreshed before the timeout expires, the join state is removed.

By default, the join state timeout is 210 seconds. You can change this timeout to allow additional time to receive the join messages. Because the messages are called join-prune messages, the name used is the `join-prune-timeout` statement.

To modify the timeout, include the `join-prune-timeout` statement:

```
user@host# set protocols pim join-prune-timeout 230
```

The join timeout value can be from 210 through 420 seconds.

SEE ALSO

| [*join-prune-timeout*](#)

Example: Enabling Join Suppression

IN THIS SECTION

- [Requirements | 322](#)
- [Overview | 323](#)
- [Configuration | 325](#)
- [Verification | 328](#)

This example describes how to enable PIM join suppression.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure PIM Sparse Mode on the interfaces. See "[Enabling PIM Sparse Mode](#)" on page 316.

Overview

IN THIS SECTION

- [Topology | 323](#)

PIM join suppression enables a router on a multiaccess network to defer sending join messages to an upstream router when it sees identical join messages on the same network. Eventually, only one router sends these join messages, and the other routers suppress identical messages. Limiting the number of join messages improves scalability and efficiency by reducing the number of messages sent to the same router.

This example includes the following statements:

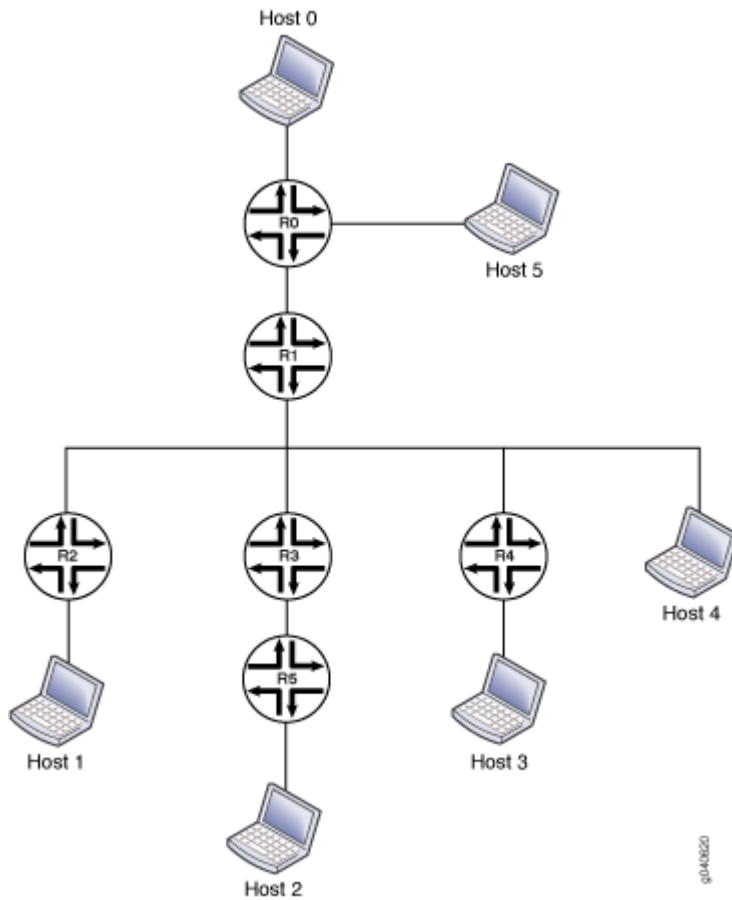
- **override-interval**—Sets the maximum time in milliseconds to delay sending override join messages. When a router sees a prune message for a join it is currently suppressing, it waits before it sends an override join message. Waiting helps avoid multiple downstream routers sending override join messages at the same time. The override interval is a random timer with a value of 0 through the maximum override value.
- **propagation-delay**—Sets a value in milliseconds for a prune pending timer, which specifies how long to wait before executing a prune on an upstream router. During this period, the router waits for any prune override join messages that might be currently suppressed. The period for the prune pending timer is the sum of the **override-interval** value and the value specified for **propagation-delay**.
- **reset-tracking-bit**—Enables PIM join suppression on each multiaccess downstream interface. This statement resets a tracking bit field (T-bit) on the LAN prune delay hello option from the default of 1 (join suppression disabled) to 0 (join suppression enabled).

When multiple identical join messages are received, a random join suppression timer is activated, with a range of 66 through 84 milliseconds. The timer is reset each time join suppression is triggered.

Topology

[Figure 42 on page 324](#) shows the topology used in this example.

Figure 42: Join Suppression



The items in the figure represent the following functions:

- Host 0 is the multicast source.
- Host 1, Host 2, Host 3, and Host 4 are receivers.
- Router R0 is the first-hop router and the RP.
- Router R1 is an upstream router.
- Routers R2, R3, R4, and R5 are downstream routers in the multicast LAN.

This example shows the configuration of the downstream devices: Routers R2, R3, R4, and R5.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 325](#)
- [Procedure | 325](#)
- [Results | 327](#)

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.

```
[edit]
set protocols pim traceoptions file pim.log
set protocols pim traceoptions file size 5m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag prune detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag register detail
set protocols pim rp static address 10.255.112.160
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set protocols pim reset-tracking-bit
set protocols pim propagation-delay 500
set protocols pim override-interval 4000
```

Procedure

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 [Junos OS CLI User Guide](#).

To configure PIM join suppression on a non-RP downstream router in the multicast LAN:

1. Configure PIM sparse mode on the interfaces.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.112.160
[edit protocols pim]
user@host# set interface all mode sparse version 2
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
```

2. Enable the join suppression timer.

```
[edit protocols pim]
user@host# set reset-tracking-bit
```

3. Configure the prune override interval value.

```
[edit protocols pim]
user@host# set override-interval 4000
```

4. Configure the propagation delay of the link.

```
[edit protocols pim]
user@host# set propagation-delay 500
```

5. (Optional) Configure PIM tracing operations.

```
[edit protocols pim]
user@host# set traceoptions file pim.log size 5m world-readable
[edit protocols pim]
user@host# set traceoptions flag join detail
[edit protocols pim]
user@host# set traceoptions flag normal detail
```



```
[edit protocols pim]
user@host# set traceoptions flag register detail
```

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

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

Results

From configuration mode, confirm your configuration by entering the `show protocols` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
pim {
  traceoptions {
    file pim.log size 5m world-readable;
    flag join detail;
    flag prune detail;
    flag normal detail;
    flag register detail;
  }
  rp {
    static {
      address 10.255.112.160;
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
  reset-tracking-bit;
  propagation-delay 500;
  override-interval 4000;
}
```

Verification

To verify the configuration, run the following commands on the upstream and downstream routers:

- `show pim join extensive`
- `show multicast route extensive`

SEE ALSO

[Example: Configuring the PIM Assert Timeout | 410](#)

[Example: Configuring PIM RPF Selection | 1167](#)

[Example: Configuring the PIM SPT Threshold Policy | 413](#)

[Enabling PIM Sparse Mode | 316](#)

[PIM Overview | 275](#)

Example: Configuring PIM Sparse Mode over an IPsec VPN

IPsec VPNs create secure point-to-point connections between sites over the Internet. The Junos OS implementation of IPsec VPNs supports multicast and unicast traffic. The following example shows how to configure PIM sparse mode for the multicast solution and how to configure IPsec to secure your traffic.


The configuration shown in this example works on the following platforms:

- M Series and T Series routers with one of the following PICs:
 - Adaptive Services (AS) PIC
 - Multiservices (MS) PIC
- JCS1200 platform with a Multiservices PIC (MS-500)

The tunnel endpoints do not need to be the same platform type. For example, the device on one end of the tunnel can be a JCS1200 router, while the device on the other end can be a standalone T Series router. The two routers that are the tunnel endpoints can be in the same autonomous system or in different autonomous systems.

In the configuration shown in this example, OSPF is configured between the tunnel endpoints. In [Figure 43 on page 329](#), the tunnel endpoints are R0 and R1. The network that contains the multicast source is connected to R0. The network that contains the multicast receivers is connected to R1. R1 serves as the statically configured rendezvous point (RP).

Figure 43: PIM Sparse Mode over an IPsec VPN



To configure PIM sparse mode with IPsec:

1. On R0, configure the incoming Gigabit Ethernet interface.

```
[edit interfaces]
user@host# set ge-0/1/1 description "incoming interface"
user@host# set ge-0/1/1 unit 0 family inet address 10.20.0.1/30
```

2. On R0, configure the outgoing Gigabit Ethernet interface.

```
[edit interfaces]
user@host# set ge-0/0/7 description "outgoing interface"
user@host# set ge-0/0/7 unit 0 family inet address 10.10.1.1/30
```

3. On R0, configure unit 0 on the sp- interface. The Junos OS uses unit 0 for service logging and other communication from the services PIC.

```
[edit interfaces]
user@host# set sp-0/2/0 unit 0 family inet
```

4. On R0, configure the logical interfaces that participate in the IPsec services. In this example, unit 1 is the inward-facing interface. Unit 1001 is the interface that faces the remote IPsec site.

```
[edit interfaces]
user@host# set sp-0/2/0 unit 1 family inet
user@host# set sp-0/2/0 unit 1 service-domain inside
user@host# set sp-0/2/0 unit 1001 family inet
user@host# set sp-0/2/0 unit 1001 service-domain outside
```

5. On R0, direct OSPF traffic into the IPsec tunnel.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface sp-0/2/0.1
user@host# set parea 0.0.0.0 interface ge-0/1/1.0 passive
user@host# set area 0.0.0.0 interface lo0.0
```

6. On R0, configure PIM sparse mode. This example uses static RP configuration. Because R0 is a non-RP router, configure the address of the RP router, which is the routable address assigned to the loopback interface on R1.

```
[edit protocols pim]
user@host# set rp static address 10.255.0.156
user@host# set interfaces sp-0/2/0.1
user@host# set interfaces ge-0/1/1.0
user@host# set interfaces lo0.0
```

7. On R0, create a rule for a bidirectional dynamic IKE security association (SA) that references the IKE policy and the IPsec policy.

```
[edit services ipsec-vpn rule ipsec_rule]
user@host# set term ipsec_dynamic then remote-gateway 10.10.1.2
user@host# set term ipsec_dynamic then dynamic ike-policy ike_policy
user@host# set term ipsec_dynamic then dynamic ipsec-policy ipsec_policy
user@host# set match-direction input
```

8. On R0, configure the IPsec proposal. This example uses the Authentication Header (AH) Protocol.

```
[edit services ipsec-vpn ipsec proposal ipsec_prop]
user@host# set protocol ah
user@host# set authentication-algorithm hmac-md5-96
```

9. On R0, define the IPsec policy.

```
[edit services ipsec-vpn ipsec policy ipsec_policy]
user@host# set perfect-forward-secrecy keys group1
user@host# set proposal ipsec_prop
```

10. On R0, configure IKE authentication and encryption details.

```
[edit services ipsec-vpn ike proposal ike_prop]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group1
user@host# set authentication-algorithm md5
user@host# set authentication-algorithm 3des-cbc
```

11. On R0, define the IKE policy.

```
[edit services ipsec-vpn ike policy ike_policy]
user@host# set proposals ike_prop
user@host# set pre-shared-key ascii-text "$ABC123"
```

12. On R0, create a service set that defines IPsec-specific information. The first command associates the IKE SA rule with IPsec. The second command defines the address of the local end of the IPsec security tunnel. The last two commands configure the logical interfaces that participate in the IPsec services. Unit 1 is for the IPsec inward-facing traffic. Unit 1001 is for the IPsec outward-facing traffic.

```
[edit services service-set ipsec_svc]
user@host# set ipsec-vpn-rules ipsec_rule
user@host# set ipsec-vpn-options local-gateway 10.10.1.1
user@host# set next-hop-service inside-service-interface sp-0/2/0.1
user@host# set next-hop-service outside-service-interface sp-0/2/0.1001
```

13. On R1, configure the incoming Gigabit Ethernet interface.

```
[edit interfaces]
user@host# set ge-2/0/1 description "incoming interface"
user@host# set ge-2/0/1 unit 0 family inet address 10.10.1.2/30
```

14. On R1, configure the outgoing Gigabit Ethernet interface.

```
[edit interfaces]
user@host# set ge-2/0/0 description "outgoing interface"
user@host# set ge-2/0/0 unit 0 family inet address 10.20.0.5/30
```

15. On R1, configure the loopback interface.

```
[edit interfaces]
user@host# set lo0.0 family inet address 10.255.0.156
```

16. On R1, configure unit 0 on the **sp-** interface. The Junos OS uses unit 0 for service logging and other communication from the services PIC.

```
[edit interfacesinterfaces]
user@host# set sp-2/1/0 unit 0 family inet
```

17. On R1, configure the logical interfaces that participate in the IPsec services. In this example, unit 1 is the inward-facing interface. Unit 1001 is the interface that faces the remote IPsec site.

```
[edit interfaces]
user@host# set sp-2/1/0 unit 1 family inet
user@host# set sp-2/1/0 unit 1 service-domain inside
user@host# set sp-2/1/0 unit 1001 family inet
user@host# set sp-2/1/0 unit 1001 service-domain outside
```

18. On R1, direct OSPF traffic into the IPsec tunnel.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface sp-2/1/0.1
user@host# set area 0.0.0.0 interface ge-2/0/0.0 passive
user@host# set area 0.0.0.0 interface lo0.0
```

19. On R1, configure PIM sparse mode. R1 is an RP router. When you configure the local RP address, use the shared address, which is the address of R1's loopback interface.

```
[edit protocols pim]
user@host# set rp local address 10.255.0.156
user@host# set interface sp-2/1/0.1
user@host# set interface ge-2/0/0.0
user@host# set interface lo0.0 family inet
```

20. On R1, create a rule for a bidirectional dynamic Internet Key Exchange (IKE) security association (SA) that references the IKE policy and the IPsec policy.

```
[edit services ipsec-vpn rule ipsec_rule]
user@host# set term ipsec_dynamic from source-address 192.168.195.34/32
user@host# set term ipsec_dynamic then remote-gateway 10.10.1.1
user@host# set term ipsec_dynamic then dynamic ike-policy ike_policy
user@host# set term ipsec_dynamic then dynamic ipsec-policy ipsec_policy
user@host# set match-direction input
```

21. On R1, define the IPsec proposal for the dynamic SA.

```
[edit services ipsec-vpn ipsec proposal ipsec_prop]
user@host# set protocol ah
user@host# set authentication-algorithm hmac-md5-96
```

22. On R1, define the IPsec policy.

```
[edit services ipsec-vpn ipsec policy ipsec_policy]
user@host# set perfect-forward-secrecy keys group1
user@host# set proposal ipsec_prop
```

23. On R1, configure IKE authentication and encryption details.

```
[edit services ipsec-vpn ike proposal ike_prop]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group1
user@host# set authentication-algorithm md5
user@host# set authentication-algorithm 3des-cbc
```

24. On R0, define the IKE policy.

```
[edit services ipsec-vpn ike policy ike_policy]
user@host# set proposal ike_prop
user@host# set pre-shared-key ascii-text "$ABC123"
```

25. On R1, create a service set that defines IPsec-specific information. The first command associates the IKE SA rule with IPsec. The second command defines the address of the local end of the IPsec

security tunnel. The last two commands configure the logical interfaces that participate in the IPsec services. Unit 1 is for the IPsec inward-facing traffic. Unit 1001 is for the IPsec outward-facing traffic.

```
[edit services service-set ipsec_svc]
user@host# set ipsec-vpn-rules ipsec_rule
user@host# set ipsec-vpn-options local-gateway 10.10.1.2
user@host# set next-hop-service inside-service-interface sp-2/1/0.1
user@host# set next-hop-service outside-service-interface sp-2/1/0.1001
```

To verify the configuration, run the following commands:

Check which RPs the various routers have learned about.

```
user@host> show pim rps extensive inet
```

Check that the IPsec SA negotiation is successful.

```
user@host> show services ipsec-vpn ipsec security-associations
```

Check that the IKE SA negotiation is successful.

```
user@host> show services ipsec-vpn ike security-associations
```

Check that traffic is traveling over the IPsec tunnel.

```
user@host> show services ipsec-vpn ipsec statistics
```

SEE ALSO

[Understanding PIM Sparse Mode | 307](#)

[Junos OS Security Services Administration Guide for Routing Devices](#)

show pim rps

[CLI Explorer](#)

[show services ipsec-vpn ipsec statistics](#)

[CLI Explorer](#)


```
show services ipsec-vpn ike security-associations
```

[CLI Explorer](#)

```
show services ipsec-vpn ipsec security-associations
```

[CLI Explorer](#)

Example: Configuring Multicast for Virtual Routers with IPv6 Interfaces

IN THIS SECTION

- [Requirements | 335](#)
- [Overview | 335](#)
- [Configuration | 336](#)
- [Verification | 341](#)

A virtual router is a type of simplified routing instance that has a single routing table. This example shows how to configure PIM in a virtual router.

Requirements

Before you begin, configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview

IN THIS SECTION

- [Topology | 336](#)

You can configure PIM for the **virtual-router** instance type as well as for the **vrf** instance type. The **virtual-router** instance type is similar to the **vrf** instance type used with Layer 3 VPNs, except that it is used for non-VPN-related applications.

The **virtual-router** instance type has no VPN routing and forwarding (VRF) import, VRF export, VRF target, or route distinguisher requirements. The **virtual-router** instance type is used for non-Layer 3 VPN situations.

When PIM is configured under the **virtual-router** instance type, the VPN configuration is not based on RFC 2547, *BGP/MPLS VPNs*, so PIM operation does not comply with the Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*. In the **virtual-router** instance type, PIM operates in a routing instance by itself, forming adjacencies with PIM neighbors over the routing instance interfaces as the other routing protocols do with neighbors in the routing instance.

This example includes the following general steps:

1. On R1, configure a virtual router instance with three interfaces (**ge-0/0/0.0**, **ge-0/1/0.0**, and **ge-0/1/1.0**).
2. Configure PIM and the RP.
3. Configure an MLD static group containing interfaces **ge-0/1/0.0** and **ge-0/1/1.0**.

After you configure this example, you should be able to send multicast traffic from R2 through **ge-0/0/0** on R1 to the static group and verify that the traffic egresses from **ge-0/1/0.0** and **ge-0/1/1.0**.




NOTE: Do not include the **group-address** statement for the **virtual-router** instance type.

Topology

Figure 44 on page 336 shows the topology for this example.

Figure 44: Virtual Router Instance with Three Interfaces



Configuration

IN THIS SECTION

- Procedure | 337
- Results | 339

Procedure

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.

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet6 address 2001:4:4:4::1/64
set interfaces ge-0/1/0 unit 0 family inet6 address 2001:24:24:24::1/64
set interfaces ge-0/1/1 unit 0 family inet6 address 2001:7:7:7::1/64
set protocols mld interface ge-0/1/0.0 static group ff0e::10
set protocols mld interface ge-0/1/1.0 static group ff0e::10
set routing-instances mvr1 instance-type virtual-router
set routing-instances mvr1 interface ge-0/0/0.0
set routing-instances mvr1 interface ge-0/1/0.0
set routing-instances mvr1 interface ge-0/1/1.0
set routing-instances mvr1 protocols pim rp local family inet6 address 2001:1:1:1::1
set routing-instances mvr1 protocols pim interface ge-0/0/0.0
set routing-instances mvr1 protocols pim interface ge-0/1/0.0
set routing-instances mvr1 protocols pim interface ge-0/1/1.0
```

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 [Junos OS CLI User Guide](#).

To configure multicast for virtual routers:

1. Configure the interfaces.

```
[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set ge-0/0/0 unit 0 family inet6 address 2001:4:4:4::1/64
[edit interfaces]
user@host# set ge-0/1/0 unit 0 family inet6 address 2001:24:24:24::1/64
[edit interfaces]
user@host# set ge-0/1/1 unit 0 family inet6 address 2001:7:7:7::1/64
```

```
[edit interfaces]
user@host# exit
```

2. Configure the routing instance type.

```
[edit]
user@host# edit routing-instances
[edit routing-instances]
user@host# set mvr1 instance-type virtual-router
```

3. Configure the interfaces in the routing instance.

```
[edit routing-instances]
user@host# set mvr1 interface ge-0/0/0
[edit routing-instances]
user@host# set mvr1 interface ge-0/1/0
[edit routing-instances]
user@host# set mvr1 interface ge-0/1/1
```

4. Configure PIM and the RP in the routing instance.

```
[edit routing-instances]
user@host# set mvr1 protocols pim rp local family inet6 address 2001:1:1:1::1
```

5. Configure PIM on the interfaces.

```
[edit routing-instances]
user@host# set mvr1 protocols pim interface ge-0/0/0
[edit routing-instances]
user@host# set mvr1 protocols pim interface ge-0/1/0
[edit routing-instances]
user@host# set mvr1 protocols pim interface ge-0/1/1
[edit routing-instances]
user@host# exit
```

6. Configure the MLD group.

```
[edit]
user@host# edit protocols mld
[edit protocols mld]
user@host# set interface ge-0/1/0.0 static group ff0e::10
[edit protocols mld]
user@host# set interface ge-0/1/1.0 static group ff0e::10
```

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

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show interfaces**, **show routing-instances**, and **show protocols** commands.

```
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet6 {
      address 2001:4:4:4::1/64;
    }
  }
}
ge-0/1/0 {
  unit 0 {
    family inet6 {
      address 2001:24:24:24::1/64;
    }
  }
}
ge-0/1/1 {
  unit 0 {
    family inet6 {
      address 2001:7:7:7::1/64;
    }
  }
}
```

```
    }  
  }
```

```
user@host# show routing-instances  
mvrfl {  
  instance-type virtual-router;  
  interface ge-0/0/0.0;  
  interface ge-0/1/0.0;  
  interface ge-0/1/1.0;  
  protocols {  
    pim {  
      rp {  
        local {  
          family inet6 {  
            address 2001:1:1:1::1;  
          }  
        }  
      }  
      interface ge-0/0/0.0;  
      interface ge-0/1/0.0;  
      interface ge-0/1/1.0;  
    }  
  }  
}
```

```
user@host# show protocols  
mld {  
  interface ge-0/1/0.0 {  
    static {  
      group ff0e::10;  
    }  
  }  
  interface ge-0/1/1.0 {  
    static {  
      group ff0e::10;  
    }  
  }  
}
```

Verification

To verify the configuration, run the following commands:

- `show mld group`
- `show mld interface`
- `show mld statistics`
- `show multicast interface`
- `show multicast route`
- `show multicast rpf`
- `show pim interfaces`
- `show pim join`
- `show pim neighbors`
- `show route forwarding-table`
- `show route instance`
- `show route table`

SEE ALSO

Configuring Virtual-Router Routing Instances in VPNs

[Junos OS VPNs Library for Routing Devices](#)

Types of VPNs

[Junos OS VPNs Library for Routing Devices](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
16.1	Starting in Junos OS Release 16.1, PIM is disabled by default. When you enable PIM, it operates in sparse mode by default.

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring PIM Bootstrap Router | 365](#)

[Configuring PIM Dense Mode | 299](#)

[Configuring a Designated Router for PIM | 425](#)

[Configuring PIM Filtering | 376](#)

[Example: Configuring Nonstop Active Routing for PIM | 515](#)

[Examples: Configuring PIM RPT and SPT Cutover | 397](#)

[Configuring PIM Sparse-Dense Mode | 304](#)

[Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol | 497](#)

[Configuring Basic PIM Settings](#)

Configuring Static RP

IN THIS SECTION

- [Understanding Static RP | 342](#)
- [Configuring Local PIM RPs | 343](#)
- [Example: Configuring PIM Sparse Mode and RP Static IP Addresses | 345](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device | 350](#)

Understanding Static RP

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the device. However, because PIM must not be configured on the network management interface, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the device must determine the IP address of the RP for the source.

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

Starting in Junos OS Release 15.2, the static configuration uses PIM version 2 by default, which is the only version supported in that release and beyond..

One common way for the device to locate RPs is by static configuration of the IP address of the RP. A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

SEE ALSO

| [Configuring Local PIM RPs | 343](#)

Configuring Local PIM RPs

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

```
[routing-instances VPN-A protocols pim]
user@host# set rp local
```

2. Configure the IP protocol family and IP address.

IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the [edit interface *interface-name*] hierarchy level and **family inet6** at the [edit protocols pim interface *interface-name*] hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
user@host# set family inet address 10.1.2.254
```

3. (IPv4 only) Configure the routing device's RP priority.



NOTE: The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set priority 5
```

4. Configure the groups for which the routing device is the RP.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set group-ranges ff70::/12
user@host# set group-ranges 224.0.0.0/4
```

5. (IPv4 only) Modify the local RP hold time.

If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not

receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set hold-time 200
```

6. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set override
```

7. Monitor the operation of PIM by running the `show pim` commands. Run `show pim ?` to display the supported commands.

SEE ALSO

[PIM Overview | 275](#)

[Understanding MLD | 59](#)

Example: Configuring PIM Sparse Mode and RP Static IP Addresses

IN THIS SECTION

- [Requirements | 346](#)
- [Overview | 346](#)
- [Configuration | 346](#)
- [Verification | 348](#)

This example shows how to configure PIM sparse mode and RP static IP addresses.

Requirements

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements.
8. Configure IGMP.

Overview

In this example, you set the interface value to **all** and disable the **ge-0/0/0** interface. Then you configure the IP address of the RP as **192.168.14.27**.

Configuration

IN THIS SECTION

- [Procedure | 347](#)

Procedure

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.

```
set protocols pim interface all
set protocols pim interface ge-0/0/0 disable
set protocols pim rp static address 192.168.14.27
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [Using the CLI Editor in Configuration Mode](#) in the [Junos OS CLI User Guide](#).

To configure PIM sparse mode and the RP static IP address:

1. Configure PIM.

```
[edit]
user@host# edit protocols pim
```

2. Set the interface value.

```
[edit protocols pim]
user@host# set pim interface all
```

3. Disable PIM on the network management interface.

```
[edit protocols pim interface]
user@host# set pim interface ge-0/0/0 unit 0 disable
```

4. Configure RP.

```
[edit]
user@host# edit protocols pim rp
```

5. Configure the IP address of the RP.

```
[edit]
user@host# set static address 192.168.14.27
```

Results

From configuration mode, confirm your configuration by entering the `show protocols` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show protocols
pim {
  rp {
    static {
      address 192.168.14.27;
    }
  }
}
interface all;
  interface ge-0/0/0.0 {
    disable;
  }
}
```

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

Verification

IN THIS SECTION

● [Verifying SAP and SDP Addresses and Ports | 349](#)

- [Verifying the IGMP Version | 349](#)
- [Verifying the PIM Mode and Interface Configuration | 349](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying SAP and SDP Addresses and Ports

Purpose

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

Action

From operational mode, enter the `show sap listen` command.

Verifying the IGMP Version

Purpose

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

Action

From operational mode, enter the `show igmp interface` command.

Verifying the PIM Mode and Interface Configuration

Purpose

Verify that PIM sparse mode is configured on all applicable interfaces.

Action

From operational mode, enter the `show pim interfaces` command.

SEE ALSO[PIM Configuration Statements](#)[Junos OS Multicast Protocols User Guide](#)[Junos OS Multicast Protocols User Guide](#)[Multicast Configuration Overview | 17](#)[Verifying a Multicast Configuration](#)**Configuring the Static PIM RP Address on the Non-RP Routing Device**

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

Starting in Junos OS Release 15.2, the default PIM version is version 2, and version 1 is not supported.

For Junos OS Release 15.1 and earlier, the default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit pim rp static address *address*]). PIM version 2 is the default for interface mode ([edit pim interface *interface-name*]). An explicitly configured PIM version will override the default setting.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

```
[routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334
```



NOTE: Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

For each static RP address, you can optionally specify the PIM version. For Junos OS Release 15.1 and earlier, the default PIM version is version 1.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
```

3. (Optional) Set the group address range.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
```

The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp static address 2001:db8:85a3::8a2e:370:7334]
user@host# set override
```

5. Monitor the operation of PIM by running the show pim commands. Run show pim ? to display the supported commands.

SEE ALSO

[PIM Overview | 275](#)

[Understanding MLD | 59](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.2	Starting in Junos OS Release 15.2, the static configuration uses PIM version 2 by default, which is the only version supported in that release and beyond.
15.2	Starting in Junos OS Release 15.2, the default PIM version is version 2, and version 1 is not supported.
15.1	For Junos OS Release 15.1 and earlier, the default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit pim rp static address <i>address</i>]). PIM version 2 is the default for interface mode ([edit pim interface <i>interface-name</i>]). An explicitly configured PIM version will override the default setting.

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring PIM Bootstrap Router | 365](#)

[Configuring a Designated Router for PIM | 425](#)

[Examples: Configuring PIM Sparse Mode | 311](#)

[Configuring Basic PIM Settings](#)

Example: Configuring Anycast RP

IN THIS SECTION

- [Understanding RP Mapping with Anycast RP | 353](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP | 353](#)
- [Example: Configuring PIM Anycast With or Without MSDP | 358](#)
- [Configuring a PIM Anycast RP Router Using Only PIM | 362](#)

Understanding RP Mapping with Anycast RP

Having a single active rendezvous point (RP) per multicast group is much the same as having a single server providing any service. All traffic converges on this single point, although other servers are sitting idle, and convergence is slow when the resource fails. In multicast specifically, there might be closer RPs on the shared tree, so the use of a single RP is suboptimal.

For the purposes of load balancing and redundancy, you can configure anycast RP. You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP fails, sources and receivers are taken to a new RP by means of unicast routing. When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

Anycast means that multiple RP routers share the same unicast IP address. Anycast addresses are advertised by the routing protocols. Packets sent to the anycast address are sent to the nearest RP with this address. Anycast addressing is a generic concept and is used in PIM sparse mode to add load balancing and service reliability to RPs.

Anycast RP is defined in RFC3446 , *Anycast RP Mechanism Using PIM and MSDP*, and can be found here: <https://www.ietf.org/rfc/rfc3446.txt> .

SEE ALSO

[Configuring Static RP | 342](#)

[Example: Configuring Multiple RPs in a Domain with Anycast RP | 353](#)

[Example: Configuring PIM Anycast With or Without MSDP | 358](#)

Example: Configuring Multiple RPs in a Domain with Anycast RP

IN THIS SECTION

- [Requirements | 354](#)
- [Overview | 354](#)
- [Configuration | 354](#)
- [Verification | 357](#)

This example shows how to configure anycast RP on each RP router in the PIM-SM domain. With this configuration you can deploy more than one RP for a single group range. This enables load balancing and redundancy.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure PIM Sparse Mode on the interfaces. See "[Enabling PIM Sparse Mode](#)" on page 316.

Overview

When you configure anycast RP, the RP routers in the PIM-SM domain use a shared address. In this example, the shared address is 10.1.1.2/32. Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources. Anycast RP also requires an RP selection method, such as static, auto-RP, or bootstrap RP. This example uses static RP and shows only one RP router configuration.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 354](#)
- [Procedure | 355](#)
- [Results | 356](#)

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.

RP Routers

```
set interfaces lo0 unit 0 family inet address 192.168.132.1/32 primary
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set protocols msdp local-address 192.168.132.1
set protocols msdp peer 192.168.12.1
set protocols pim rp local address 10.1.1.2
set routing-options router-id 192.168.132.1
```

Non-RP Routers

```
set protocols pim rp static address 10.1.1.2
```

Procedure

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 [Junos OS CLI User Guide](#).

To configure anycast RP:

1. On each RP router in the domain, configure the shared anycast address on the router's loopback address.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.1.1.2/32
```

2. On each RP router in the domain, make sure that the router's regular loopback address is the primary address for the interface, and set the router ID.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 192.168.132.1/32 primary
[edit routing-options]
user@host# set router-id 192.168.132.1
```

3. On each RP router in the domain, configure the local RP address, using the shared address.

```
[edit protocols pim]
user@host# set rp local address 10.1.1.2
```

4. On each RP router in the domain, create MSDP sessions to the other RPs in the domain.

```
[edit protocols msdp]
user@host# set local-address 192.168.132.1
user@host# set peer 192.168.12.1
```

5. On each non-RP router in the domain, configure a static RP address using the shared address.

```
[edit protocols pim]
user@host# set rp static address 10.1.1.2
```

6. If you are done configuring the devices, commit the configuration.

```
user@host# 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@host# show interfaces
lo0 {
  unit 0 {
    family inet {
      address 192.168.132.1/32 {
        primary;
      }
      address 10.1.1.2/32;
    }
  }
}
```

On the RP routers:

```
user@host# show protocols
msdp {
  local-address 192.168.132.1;
  peer 192.168.12.1;
}
pim {
  rp {
    local {
      address 10.1.1.2;
    }
  }
}
```

On the non-RP routers:

```
user@host# show protocols
pim {
  rp {
    static {
      address 10.1.1.2;
    }
  }
}
```

```
user@host# show routing-options
router-id 192.168.132.1;
```

Verification

To verify the configuration, run the **show pim rps extensive inet** command.

SEE ALSO

[Example: Configuring PIM Anycast With or Without MSDP | 358](#)

[Understanding PIM Sparse Mode | 307](#)

[Understanding RP Mapping with Anycast RP | 353](#)

Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

Starting in Junos OS Release 16.1, all systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the `[edit protocols pim rp static address address]` hierarchy level). However, PIMv2 is the default for interface mode (at the `[edit protocols pim interface interface-name]` hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the `lo0` loopback interface, which is always up. Include the `address` statement and specify the unique and routable router ID and the RP address at the `[edit interfaces lo0 unit 0 family inet]` hierarchy level. In this example, the router ID is `198.51.100.254` and the shared RP address is `198.51.100.253`. Include the `primary` statement for the first address. Including the `primary` statement selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.51.100.254/32;
        primary;
        address 198.51.100.253/32;
      }
    }
  }
}

```



```

    }
  }
}

```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```

protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.51.100.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}

```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```

protocols {
  msdp {
    peer 198.51.100.250 {
      local-address address 198.51.100.254;
    }
  }
}

```

```

    }
}

```



NOTE: If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step "4" on page 360 and Step "5" on page 360. Otherwise, go to Step "6" on page 361.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.51.100.254** and the shared RP address is **198.51.100.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.51.100.254/32 {
          primary;
        }
        address 198.51.100.253/32;
      }
    }
  }
}

```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```

protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.51.100.253;
          anycast-pim {
            rp-set {
              address 198.51.100.240;
              address 198.51.100.241 forward-msdp-sa;
            }
            local-address 198.51.100.254; #If not configured, use lo0 primary
          }
        }
      }
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```

protocols {
  pim {
    rp {
      static {
        address 198.51.100.253 {
          version 2;
        }
      }
    }
  }
}

```

```

    }
  }
}

```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```

protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Configuring a PIM Anycast RP Router Using Only PIM

In this example, configure an RP using the **lo0** loopback interface, which is always up. Use the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this case, the router ID is 198.51.100.254/32 and the shared RP address is 198.51.100.253/32. Add the flag statement **primary** to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.51.100.254/32 {
          primary;
        }
        address 198.51.100.253/32;
      }
    }
  }
}

```

```

    }
}

```

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse**, and include the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

```

protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.51.100.253;
          anycast-pim {
            rp-set {
              address 198.51.100.240;
              address 198.51.100.241 forward-msdp-sa;
            }
            local-address 198.51.100.254; #If not configured, use lo0 primary
          }
        }
      }
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}
}

```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

Use the `show pim join` and `show pim rps` commands to confirm:

```

user@device> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: <group IP>
  Source: <source IP>
  Flags: sparse,spt
  Upstream interface: lo0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout: 357
  Uptime: 00:00:03
  Downstream neighbors:
  Number of downstream interfaces: 0
  Number of downstream neighbors: 0

```

```

user@device> show pim rps extensive
Instance: PIM.master

address-family INET

RP: 198.51.100.253
Learned via: static configuration
Mode: Sparse
Time Active: 00:24:14
Holdtime: 150
Device Index: 150
Subunit: 32700
Interface:
Static RP Override: Off
Group Ranges:
  <group IP>
Register State for RP:

```

Group	Source	FirstHop	RP Address	State	Timeout
<group IP>	<source IP>	<IP address>	198.51.100.253	Receive	170

```

Anycast PIM local address used: <IP address>

```

SEE ALSO

[JTAC Certified Step-by-Step Troubleshooting: Junos OS Multicast](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
16.1	Starting in Junos OS Release 16.1, all systems on a subnet must run the same version of PIM.

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring PIM Bootstrap Router | 365](#)

[Configuring a Designated Router for PIM | 425](#)

[Examples: Configuring PIM Sparse Mode | 311](#)

[Configuring Basic PIM Settings](#)

Configuring PIM Bootstrap Router**IN THIS SECTION**

- [Understanding the PIM Bootstrap Router | 365](#)
- [Configuring PIM Bootstrap Properties for IPv4 | 366](#)
- [Configuring PIM Bootstrap Properties for IPv4 or IPv6 | 367](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain | 369](#)
- [Example: Configuring PIM BSR Filters | 370](#)

Understanding the PIM Bootstrap Router

To determine which router is the rendezvous point (RP), all routers within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group of routers that all share the same RP router. The domain bootstrap router initiates bootstrap messages, which are sent hop by hop within the

domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

Configuring PIM Bootstrap Properties for IPv4

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The bootstrap router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.



NOTE: For legacy configuration purposes, there are two sections that describe the configuration of bootstrap routers: one section for both IPv4 and IPv6, and this section, which is for IPv4 only. The method described in Configuring PIM Bootstrap Properties for IPv4 or IPv6 is recommended. A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is “duplicate IPv4 bootstrap configuration.”

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. A priority of 0 disables the function for IPv4 and does not cause the routing device to send bootstrap router packets with a 0 in the priority field. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.

```
[edit protocols pim rp]
user@host# set bootstrap-priority 3
```

2. (Optional) Create import and export policies to control the flow of IPv4 bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routing devices in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The `bootstrap-import`

statement prevents messages from being imported into the RP. The bootstrap-export statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap-import pim-bootstrap-import
user@host# set bootstrap-export pim-bootstrap-export
```

3. Configure the policies.

```
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
[edit policy-options policy-statement pim-bootstrap-export]
user@host# set from interface se-0/0/0
user@host# set then reject
```

4. Monitor the operation of PIM bootstrap routing devices by running the show pim bootstrap command.

SEE ALSO

[Understanding PIM Sparse Mode | 307](#)

show pim bootstrap

[CLI Explorer](#)

Configuring PIM Bootstrap Properties for IPv4 or IPv6

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The bootstrap router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.



NOTE: For legacy configuration purposes, there are two sections that describe the configuration of bootstrap routers: one section for IPv4 only, and this section, which is for both IPv4 and IPv6. The method described in this section is recommended. A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is “duplicate IPv4 bootstrap configuration.”

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are

configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.



NOTE: In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the configuration shown here, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the `bootstrap` statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routing devices in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The `import` statement prevents messages from being imported into the RP. The `export` statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
```

3. Configure the policies.

```
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
```

```

user@host# set then reject
user@host# exit
user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject

```

4. Monitor the operation of PIM bootstrap routing devices by running the `show pim bootstrap` command.

SEE ALSO

[Understanding PIM Sparse Mode | 307](#)

show pim bootstrap

[CLI Explorer](#)

Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the **from interface so-0-1/0 then reject** policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```

protocols {
  pim {
    rp {
      bootstrap {
        family inet {
          priority 1;
          import pim-import;
          export pim-export;
        }
        family inet6 {
          priority 1;
          import pim-import;
          export pim-export;
        }
      }
    }
  }
}
policy-options {
  policy-statement pim-import {
    from interface so-0/1/0;
    then reject;
  }
}

```

```
    }  
    policy-statement pim-export {  
      to interface so-0/1/0;  
      then reject;  
    }  
  }  
}
```

Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {  
  pim {  
    rp {  
      bootstrap-import no-bsr;  
      bootstrap-export no-bsr;  
    }  
  }  
}  
policy-options {  
  policy-statement no-bsr {  
    then reject;  
  }  
}
```

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring a Designated Router for PIM | 425](#)

[Examples: Configuring PIM Sparse Mode | 311](#)

[Configuring Basic PIM Settings](#)

Understanding PIM Auto-RP

You can configure a more dynamic way of assigning rendezvous points (RPs) in a multicast network by means of auto-RP. When you configure auto-RP for a router, the router learns the address of the RP in the network automatically and has the added advantage of operating in PIM version 1 and version 2.

Although auto-RP is a nonstandard (non-RFC-based) function that typically uses dense mode PIM to advertise control traffic, it provides an important failover advantage that simple static RP assignment does not. You can configure multiple routers as RP candidates. If the elected RP fails, one of the other preconfigured routers takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

RELATED DOCUMENTATION

| [Configuring PIM Auto-RP](#)

Configuring All PIM Anycast Non-RP Routers

Use the `mode` statement at the `[edit protocols pim rp interface all]` hierarchy level to specify sparse mode on all interfaces. Then add the `version` statement at the `[edit protocols pim rp interface all mode]` to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the `fxp0.0` management interface by adding the `disable` statement for that interface.

```
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring a PIM Anycast RP Router with MSDP

Add the `address` statement at the `[edit protocols pim rp local]` hierarchy level to specify the RP address (the same address as the secondary `lo0` interface).

For all interfaces, use the `mode` statement to set the mode to **sparse** and the `version` statement to specify PIM version 2 at the `[edit protocols pim rp local interface all]` hierarchy level. When configuring all interfaces, exclude the `fxp0.0` management interface by adding the `disable` statement for that interface.

```
protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.51.100.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

To configure MSDP peering, add the `peer` statement to configure the address of the MSDP peer at the `[edit protocols msdp]` hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the `local-address` statement at the `[edit protocols msdp peer]` hierarchy level.

```
protocols {
  msdp {
    peer 198.51.100.250 {
      local-address 198.51.100.254;
    }
  }
}
```

Configuring Embedded RP

IN THIS SECTION

- [Understanding Embedded RP for IPv6 Multicast | 373](#)
- [Configuring PIM Embedded RP for IPv6 | 375](#)

Understanding Embedded RP for IPv6 Multicast

Global IPv6 multicast between routing domains has been possible only with source-specific multicast (SSM) because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks, this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. This feature embeds the RP address in an IPv6 multicast address.

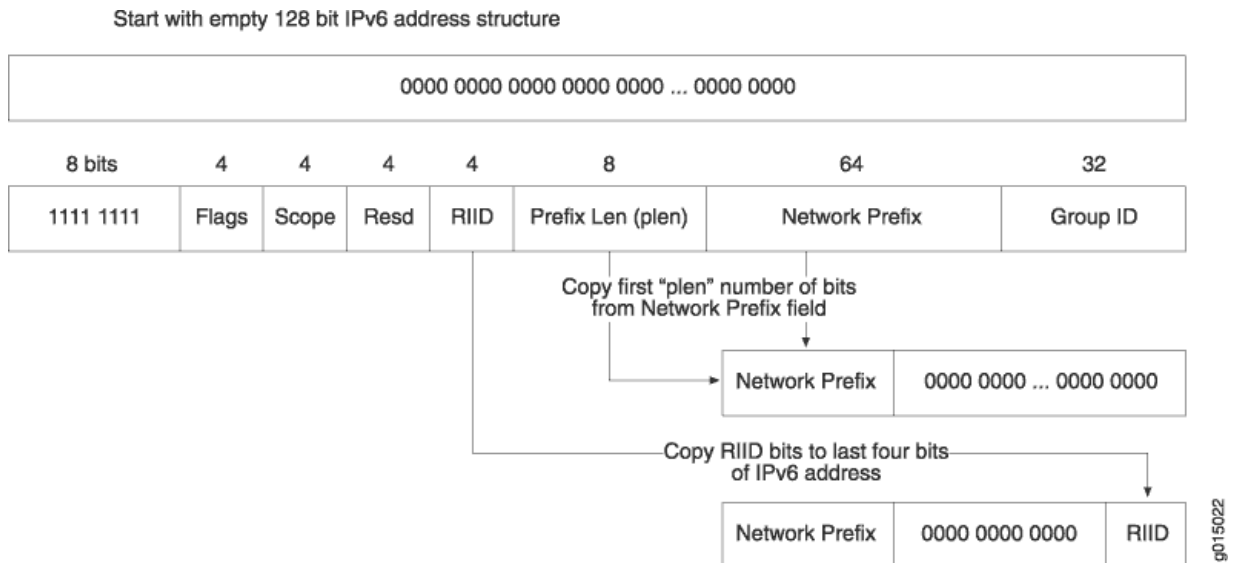
All IPv6 multicast addresses begin with 8 1-bits (1111 1111) followed by a 4-bit flag field normally set to 0011. The flag field is set to 0111 when embedded RP is used. Then the low-order bits of the normally reserved field in the IPv6 multicast address carry the 4-bit RP interface identifier (RIID).

When the IPv6 address of the RP is embedded in a unicast-prefix-based any-source multicast (ASM) address, all of the following conditions must be true:

- The address must be an IPv6 multicast address and have 0111 in the flags field (that is, the address is part of the prefix FF70::/12).
- The 8-bit prefix length (plen) field must not be all 0. An all 0 plen field implies that SSM is in use.
- The 8-bit prefix length field value must not be greater than 64, which is the length of the network prefix field in unicast-prefix-based ASM addresses.

The routing platform derives the value of the interdomain RP by copying the prefix length field number of bits from the 64-bit network prefix field in the received IPv6 multicast address to an empty 128-bit IPv6 address structure and copying the last bits from the 4-bit RIID. For example, if the prefix length field bits have the value 32, then the routing platform copies the first 32 bits of the IPv6 multicast address network prefix field to an all-0 IPv6 address and appends the last four bits determined by the RIID. See [Figure 45 on page 374](#) for an illustration of this process.

Figure 45: Extracting the Embedded RP IPv6 Address



For example, the administrator of IPv6 network 2001:DB8::/32 sets up an RP for the 2001:DB8:BEEF:FEED::/96 subnet. In that case, the received embedded RP IPv6 ASM address has the form:

```
FF70:y40:2001:DB8:BEEF:FEED::/96
```

and the derived RP IPv6 address has the form:

```
2001:DB8:BEEF:FEED::y
```

where **y** is the RIID (**y** cannot be 0).

When configured, the routing platform checks for embedded RP information in every PIM join request received for IPv6. The use of embedded RP does not change the processing of IPv6 multicast and RPs in any way, except that the embedded RP address is used if available and selected for use. There is no need to specify the IPv6 address family for embedded RP configuration because the information can be used only if IPv6 multicast is properly configured on the routing platform.

The following receive events trigger extraction of an IPv6 embedded RP address on the routing platform:

- Multicast Listener Discovery (MLD) report for an embedded RP multicast group address
- PIM join message with an embedded RP multicast group address
- Static embedded RP multicast group address associated with an interface

- Packets sent to an embedded RP multicast group address received on the DR

The embedded RP node discovered through these events is added if it does not already exist on the routing platform. The routing platform chooses the embedded RP as the RP for a multicast group before choosing an RP learned through BSRs or a statically configured RP. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.

Configuring PIM Embedded RP for IPv6

You configure embedded RP to allow multidomain IPv6 multicast networks to find RPs in other routing domains. Embedded RP embeds an RP address inside PIM join messages and other types of messages sent between routing domains. Global IPv6 multicast between routing domains has been possible only with source-specific multicast (SSM) because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks, this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. Thus, embedded RP enables you can deploy IPv6 with any-source multicast (ASM).

Embedded RP is disabled by default.

When you configure embedded RP for IPv6, embedded RPs are preferred to RPs discovered by IPv6 any other way. You configure embedded RP independent of any other IPv6 multicast properties. This feature is applied only when IPv6 multicast is properly configured.

You can configure embedded RP globally or for a routing instance. This example shows the routing instance configuration.

To configure embedded RP for IPv6 PIM sparse mode:

1. Define which multicast addresses or prefixes can embed RP address information. If messages within a group range contain embedded RP information and the group range is not configured, the embedded RP in that group range is ignored. Any valid unicast-prefix-based ASM address can be used as a group range. The default group range is FF70::/12 to FFF0::/12. Messages with embedded RP information that do not match any configured group ranges are treated as normal multicast addresses.

```
[edit routing-instances vpn-A protocols pim rp embedded-rp]
user@host# set group-ranges fec0::/10
```

If the derived RP address is not a valid IPv6 unicast address, it is treated as any other multicast group address and is not used for RP information. Verification fails if the extracted RP address is a local interface, unless the routing device is configured as an RP and the extracted RP address matches the configured RP address. Then the local RP determines whether it is configured to act as an RP for the embedded RP multicast address.

2. Limit the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.

```
[edit routing-instances vpn-A protocols pim rp]  
user@host# set maximum-rps 50
```

3. Monitor the operation by running the `show pim rps` and `show pim statistics` commands.

SEE ALSO

show pim rps

[CLI Explorer](#)

show pim statistics

[CLI Explorer](#)

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring PIM Bootstrap Router | 365](#)

[Configuring a Designated Router for PIM | 425](#)

[Examples: Configuring PIM Sparse Mode | 311](#)

[Configuring Basic PIM Settings](#)

Configuring PIM Filtering

IN THIS SECTION

- [Understanding Multicast Message Filters | 377](#)
- [Filtering MAC Addresses | 378](#)
- [Filtering RP and DR Register Messages | 378](#)
- [Filtering MSDP SA Messages | 379](#)
- [Configuring Interface-Level PIM Neighbor Policies | 380](#)
- [Filtering Outgoing PIM Join Messages | 381](#)

- [Example: Stopping Outgoing PIM Register Messages on a Designated Router | 382](#)
- [Filtering Incoming PIM Join Messages | 387](#)
- [Example: Rejecting Incoming PIM Register Messages on RP Routers | 389](#)
- [Configuring Register Message Filters on a PIM RP and DR | 394](#)

Understanding Multicast Message Filters

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.



NOTE: If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.



NOTE: If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

SEE ALSO

[Filtering MAC Addresses | 378](#)

Filtering MAC Addresses

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF (multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

Filtering RP and DR Register Messages

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

- Group multicast address

- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The `route-filter` statement is useful for multicast group address filtering, and the `source-address-filter` statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.

SEE ALSO

[Understanding RP Mapping with Anycast RP | 353](#)

Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network, you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.



NOTE: When you apply firewall filters, firewall action modifiers, such as **log**, **sample**, and **count**, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

SEE ALSO

[Filtering Incoming PIM Join Messages | 387](#)

Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

```
[edit policy-options]
user@host# set prefix-list nbrGroup 1 20.20.20.1/32
user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
user@host# set policy-statement nbr-policy then reject
```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set neighbor-policy nbr-policy
```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the `show pim statistics` command.

SEE ALSO

show pim statistics

Filtering Outgoing PIM Join Messages

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

```
user@host# set policy-options policy-statement block-groups term t1 from route-filter
224.0.1.2/32 exact
user@host# set policy-options policy-statement block-groups term t1 from route-filter
225.1.1.1/32 exact
user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept
```

2. Verify the policy configuration by running the `show policy-options` command.

```
user@host# show policy-options
policy-statement block-groups {
  term t1 {
    from {
      route-filter 224.0.1.2/32 exact;
      route-filter 225.1.1.1/32 exact;
    }
    then reject;
  }
  term last {
    then accept;
  }
}
```

3. Apply the PIM join and prune message filter.

```
user@host> set protocols pim export block-groups
```

4. After the configuration is committed, use the `show pim statistics` command to verify that outgoing PIM join and prune messages are being filtered.

```
user@host> show pim statistics | grep filtered
RP Filtered Source          0
Rx Joins/Prunes filtered    0
Tx Joins/Prunes filtered    254
```

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

SEE ALSO

| [Filtering Incoming PIM Join Messages | 387](#)

Example: Stopping Outgoing PIM Register Messages on a Designated Router

IN THIS SECTION

- [Requirements | 382](#)
- [Overview | 383](#)
- [Configuration | 383](#)
- [Verification | 386](#)

This example shows how to stop outgoing PIM register messages on a designated router.

Requirements

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.

2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements.
8. Configure IGMP.
9. Configure the PIM static RP.
10. Filter PIM register messages from unauthorized groups and sources. See ["Example: Rejecting Incoming PIM Register Messages on RP Routers"](#) on page 389.

Overview

In this example, you configure the group address as **224.2.2.2/32** and the source address in the group as **20.20.20.1/32**. You set the match action to not send PIM register messages for the group and source address. Then you configure the policy on the designated router to **stop-pim-register-msg-dr**.

Configuration

IN THIS SECTION

- [Procedure | 384](#)

Procedure

CLI Quick Configuration

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

```
set policy-options policy-statement stop-pim-register-msg-dr from route-filter 224.2.2.2/32 exact
set policy-options policy-statement stop-pim-register-msg-dr from source-address-filter 20.20.20.1/32 exact
set policy-options policy-statement stop-pim-register-msg-dr then reject
set protocols pim rp dr-register-policy stop-pim-register-msg-dr
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [Using the CLI Editor in Configuration Mode](#) in the [Junos OS CLI User Guide](#).

To stop outgoing PIM register messages on a designated router:

1. Configure the policy options.

```
[edit]
user@host# edit policy-options
```

2. Set the group address.

```
[edit policy-options]
user@host# set policy statement stop-pim-register-msg-dr from route-filter 224.2.2.2/32 exact
```

3. Set the source address.

```
[edit policy-options]
user@host# set policy statement stop-pim-register-msg-dr from source-address-filter 20.20.20.1/32 exact
```

4. Set the match action.

```
[edit policy-options]
user@host# set policy statement stop-pim-register-msg-dr then reject
```

5. Assign the policy.

```
[edit]
user@host# set dr-register-policy stop-pim-register-msg-dr
```

Results

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

```
[edit]
user@host# show policy-options
policy-statement stop-pim-register-msg-dr {
  from {
    route-filter 224.2.2.2/32 exact;
    source-address-filter 20.20.20.1/32 exact;
  }
  then reject;
}
[edit]
user@host# show protocols
pim {
  rp {
    dr-register-policy stop-pim-register-msg-dr;
  }
}
```

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

Verification

IN THIS SECTION

- [Verifying SAP and SDP Addresses and Ports | 386](#)
- [Verifying the IGMP Version | 386](#)
- [Verifying the PIM Mode and Interface Configuration | 386](#)
- [Verifying the PIM RP Configuration | 387](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying SAP and SDP Addresses and Ports

Purpose

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

Action

From operational mode, enter the `show sap listen` command.

Verifying the IGMP Version

Purpose

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

Action

From operational mode, enter the `show igmp interface` command.

Verifying the PIM Mode and Interface Configuration

Purpose

Verify that PIM sparse mode is configured on all applicable interfaces.

Action

From operational mode, enter the `show pim interfaces` command.

Verifying the PIM RP Configuration**Purpose**

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

Action

From operational mode, enter the `show pim rps` command.

SEE ALSO

[Multicast Configuration Overview | 17](#)

Filtering Incoming PIM Join Messages

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See [Table 13 on page 387](#) for a list of match conditions.

Table 13: PIM Join Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address

Table 13: PIM Join Filter Match Conditions (Continued)

Match Condition	Matches On
neighbor	Neighbor address (the source address in the IP header of the join and prune message)
route-filter	Multicast group address embedded in the join and prune message
source-address-filter	Multicast source address embedded in the join and prune message

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. The **bad-groups** filter prevents (*,G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

```
[edit policy-statement pim-join-filter term bad-groups]
user@host# set from route-filter 224.0.1.2/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject
```

```
[edit policy-statement pim-join-filter term bad-sources]
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

```
[edit policy-statement pim-join-filter term last]
user@host# set then accept
```

2. Apply one or more policies to routes being imported into the routing table from PIM.

```
[edit protocols pim]
user@host# set import pim-join-filter
```

3. Verify the configuration by checking the output of the `show pim join` and `show policy` commands.

SEE ALSO

[Understanding Multicast Administrative Scoping | 1273](#)

[Filtering Outgoing PIM Join Messages | 381](#)

`show pim join`

[CLI Explorer](#)

`show policy`

[CLI Explorer](#)

Example: Rejecting Incoming PIM Register Messages on RP Routers

IN THIS SECTION

- [Requirements | 389](#)
- [Overview | 390](#)
- [Configuration | 390](#)
- [Verification | 393](#)

This example shows how to reject incoming PIM register messages on RP routers.

Requirements

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.

3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See ["Configuring the Session Announcement Protocol" on page 573](#).
8. Configure IGMP. See ["Configuring IGMP" on page 25](#).
9. Configure the PIM static RP. See ["Configuring Static RP" on page 342](#).

Overview

In this example, you configure the group address as **224.1.1.1/32** and the source address in the group as **10.10.10.1/32**. You set the match action to reject PIM register messages and assign `reject-pim-register-msg-rp` as the policy on the RP.

Configuration

IN THIS SECTION

- [Procedure | 390](#)

Procedure

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.

```
set policy-options policy-statement reject-pim-register-msg-rp from route-filter 224.1.1.1/32
exact
set policy-options policy-statement reject-pim-register-msg-rp from source-address-filter
```



```
10.10.10.1/32 exact
set policy-options policy-statement reject-pim-register-msg-rp then reject
set protocols pim rp rp-register-policy reject-pim-register-msg-rp
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [Using the CLI Editor in Configuration Mode](#) in the [Junos OS CLI User Guide](#).

To reject the incoming PIM register messages on an RP router:

1. Configure the policy options.

```
[edit]
user@host# edit policy-options
```

2. Set the group address.

```
[edit policy-options]
user@host# set policy statement reject-pim-register-msg-rp from route-filter 224.1.1.1/32
exact
```

3. Set the source address.

```
[edit policy-options]
user@host# set policy statement reject-pim-register-msg-rp from source-address-filter
10.10.10.1/32 exact
```

4. Set the match action.

```
[edit policy-options]
user@host# set policy statement reject-pim-register-msg-rp then reject
```

5. Configure the protocol.

```
[edit]
user@host# edit protocols pim rp
```

6. Assign the policy.

```
[edit]
user@host# set rp-register-policy reject-pim-register-msg-rp
```

Results

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

```
[edit]
user@host# show policy-options
policy-statement reject-pim-register-msg-rp {
  from {
    route-filter 224.1.1.1/32 exact;
    source-address-filter 10.10.10.1/32 exact;
  }
  then reject;
}
[edit]
user@host# show protocols pim
rp {
  rp-register-policy reject-pim-register-msg-rp;
}
```

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

Verification

IN THIS SECTION

- [Verifying SAP and SDP Addresses and Ports | 393](#)
- [Verifying the IGMP Version | 393](#)
- [Verifying the PIM Mode and Interface Configuration | 393](#)
- [Verifying the PIM Register Messages | 394](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying SAP and SDP Addresses and Ports

Purpose

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

Action

From operational mode, enter the `show sap listen` command.

Verifying the IGMP Version

Purpose

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

Action

From operational mode, enter the `show igmp interface` command.

Verifying the PIM Mode and Interface Configuration

Purpose

Verify that PIM sparse mode is configured on all applicable interfaces.

Action

From operational mode, enter the `show pim interfaces` command.

Verifying the PIM Register Messages

Purpose

Verify whether the rejected policy on the RP router is enabled.

Action

From configuration mode, enter the `show policy-options` and `show protocols pim` command.

SEE ALSO

[Example: Stopping Outgoing PIM Register Messages on a Designated Router | 382](#)

[Multicast Configuration Overview | 17](#)

[Verifying a Multicast Configuration](#)

Configuring Register Message Filters on a PIM RP and DR

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP router.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

```
[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy [ reject_224_1_1_1 | accept_224_1_1_5 ]
user@host# set local address 10.10.10.5
```

To monitor the operation of the filters, run the `show pim statistics` command. The command output contains the following fields related to filtering:

- RP Filtered Source
- Rx Joins/Prunes filtered
- Tx Joins/Prunes filtered
- Rx Register msgs filtering drop
- Tx Register msgs filtering drop

SEE ALSO

[PIM Sparse Mode Source Registration](#) | 400

[Filtering RP and DR Register Messages | 378](#)

show pim statistics

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring PIM Bootstrap Router | 365](#)

[Configuring PIM Dense Mode | 299](#)

[Configuring a Designated Router for PIM | 425](#)

[Example: Configuring Nonstop Active Routing for PIM | 515](#)

[Examples: Configuring PIM RPT and SPT Cutover | 397](#)

[Configuring PIM Sparse-Dense Mode | 304](#)

[Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol | 497](#)

[Configuring Basic PIM Settings](#)

Examples: Configuring PIM RPT and SPT Cutover

IN THIS SECTION

- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees | 398](#)
- [Building an RPT Between the RP and Receivers | 399](#)
- [PIM Sparse Mode Source Registration | 400](#)
- [Multicast Shortest-Path Tree | 403](#)
- [SPT Cutover | 404](#)
- [SPT Cutover Control | 409](#)
- [Example: Configuring the PIM Assert Timeout | 410](#)
- [Example: Configuring the PIM SPT Threshold Policy | 413](#)

Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees

In a shared tree, the root of the distribution tree is a router, not a host, and is located somewhere in the core of the network. In the primary sparse mode multicast routing protocol, Protocol Independent Multicast sparse mode (PIM SM), the core router at the root of the shared tree is the rendezvous point (RP). Packets from the upstream source and join messages from the downstream routers “rendezvous” at this core router.

In the RP model, other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups.

The RP model shifts the burden of finding sources of multicast content from each router (the (S,G) notation) to the network (the (*,G) notation knows only the RP). Exactly how the RP finds the unicast IP address of the source varies, but there must be some method to determine the proper source for multicast content for a particular group.

Consider a set of multicast routers without any active multicast traffic for a certain group. When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the distribution tree for that group back to the RP, not to the actual source of the content.

To join the shared tree, or () as it is called in PIM sparse mode, the router must do the following:

- Determine the IP address of the RP for that group. Determining the address can be as simple as static configuration in the router, or as complex as a set of nested protocols.
- Build the shared tree for that group. The router executes an RPF check on the RP address in its routing table, which produces the interface closest to the RP. The router now detects that multicast packets from this RP for this group need to flow into the router on this RPF interface.
- Send a join message out on this interface using the proper multicast protocol (probably PIM sparse mode) to inform the upstream router that it wants to join the shared tree for that group. This message is a (*,G) join message because S is not known. Only the RP is known, and the RP is not actually the source of the multicast packets. The router receiving the (*,G) join message adds the interface on which the message was received to its outgoing interface list (OIL) for the group and also performs an RPF check on the RP address. The upstream router then sends a (*,G) join message out from the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating join messages from the RPF interface, building the shared tree as it goes. The process stops when the join message reaches one of the following:

- The RP for the group that is being joined

- A router along the RPT that already has a multicast forwarding state for the group that is being joined

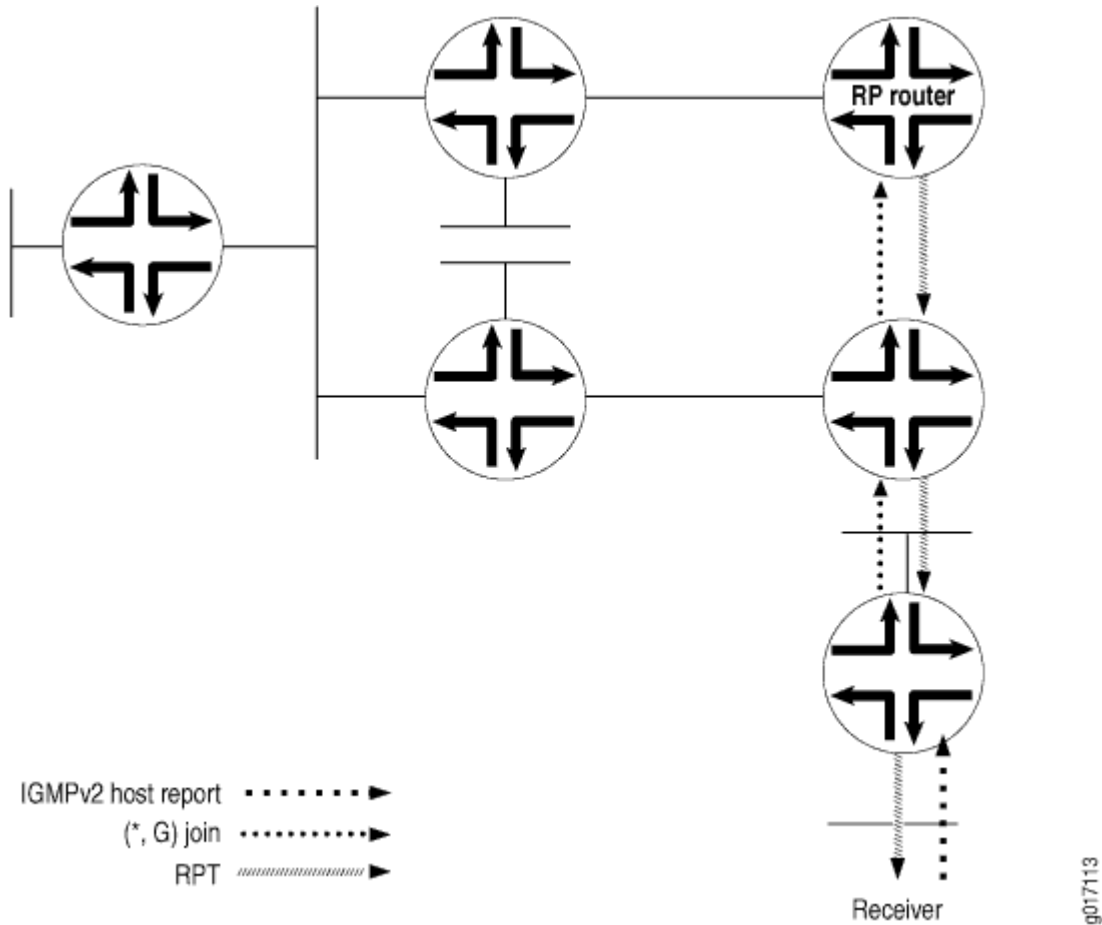
In either case, the branch is created, and packets can flow from the source to the RP and from the RP to the receiver. Note that there is no guarantee that the shared tree (RPT) is the shortest path tree to the source. Most likely it is not. However, there are ways to “migrate” a shared tree to an SPT once the flow of packets begins. In other words, the forwarding state can transition from (*,G) to (S,G). The formation of both types of tree depends heavily on the operation of the RPF check and the RPF table. For more information about the RPF table, see ["Understanding Multicast Reverse Path Forwarding" on page 1149](#).

Building an RPT Between the RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see [Figure 46 on page 400](#)). The RPT is built by means of a PIM join message from a receiver's DR:

1. A receiver sends a request to join group (G) in an Internet Group Management Protocol (IGMP) host membership report. A PIM sparse-mode router, the receiver's DR, receives the report on a directly attached subnet and creates an RPT branch for the multicast group of interest.
2. The receiver's DR sends a PIM join message to its RPF neighbor, the next-hop address in the RPF table, or the unicast routing table.
3. The PIM join message travels up the tree and is multicast to the ALL-PIM-ROUTERS group (224.0.0.13). Each router in the tree finds its RPF neighbor by using either the RPF table or the unicast routing table. This is done until the message reaches the RP and forms the RPT. Routers along the path set up the multicast forwarding state to forward requested multicast traffic back down the RPT to the receiver.

Figure 46: Building an RPT Between the RP and the Receiver




PIM Sparse Mode Source Registration

The RPT is a unidirectional tree, permitting traffic to flow down from the RP to the receiver in one direction. For multicast traffic to reach the receiver from the source, another branch of the distribution tree, called the shortest-path tree, needs to be built from the source's DR to the RP.

The shortest-path tree is created in the following way:


1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see [Figure 47 on page 401](#)).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 47: PIM Register Message and PIM Join Message Exchanged



3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see [Figure 48 on page 402](#)).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.


Figure 48: Traffic Sent from the Source to the RP Router



g017115

5. The RP router sends the multicast traffic down the RPT toward the receiver (see [Figure 49 on page 403](#)).

Figure 49: Traffic Sent from the RP Router Toward the Receiver



Multicast Shortest-Path Tree

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT) as well. Consider a set of multicast routers without any active multicast traffic for a certain group (that is, they have no multicast forwarding state for that group). When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the tree for that group.

To join the distribution tree, the router determines the unicast IP address of the source for that group. This address can be a simple static configuration on the router, or as complex as a set of protocols.

To build the SPT for that group, the router executes an a reverse path forwarding (RPF) check on the source address in its routing table. The RPF check produces the interface closest to the source, which is where multicast packets from this source for this group need to flow into the router.

The router next sends a join message out on this interface using the proper multicast protocol to inform the upstream router that it wants to join the distribution tree for that group. This message is an (S,G) join message because both S and G are known. The router receiving the (S,G) join message adds the interface on which the message was received to its output interface list (OIL) for the group and also performs an

RPF check on the source address. The upstream router then sends an (S,G) join message out on the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating joins out on the RPF interface, building the SPT as it goes. The process stops when the join message does one of two things:

- Reaches the router directly connected to the host that is the source.
- Reaches a router that already has multicast forwarding state for this source-group pair.

In either case, the branch is created, each of the routers has multicast forwarding state for the source-group pair, and packets can flow down the distribution tree from source to receiver. The RPF check at each router makes sure that the tree is an SPT.


SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point. For more information about RPs, see "[Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees](#)" on page 398.

SPT Cutover

Instead of continuing to use the SPT to the RP and the RPT toward the receiver, a direct SPT is created between the source and the receiver in the following way:


1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see [Figure 50 on page 405](#)).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

Figure 50: Receiver DR Sends a PIM Join Message to the Source




4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see [Figure 51 on page 406](#)).

Figure 51: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router



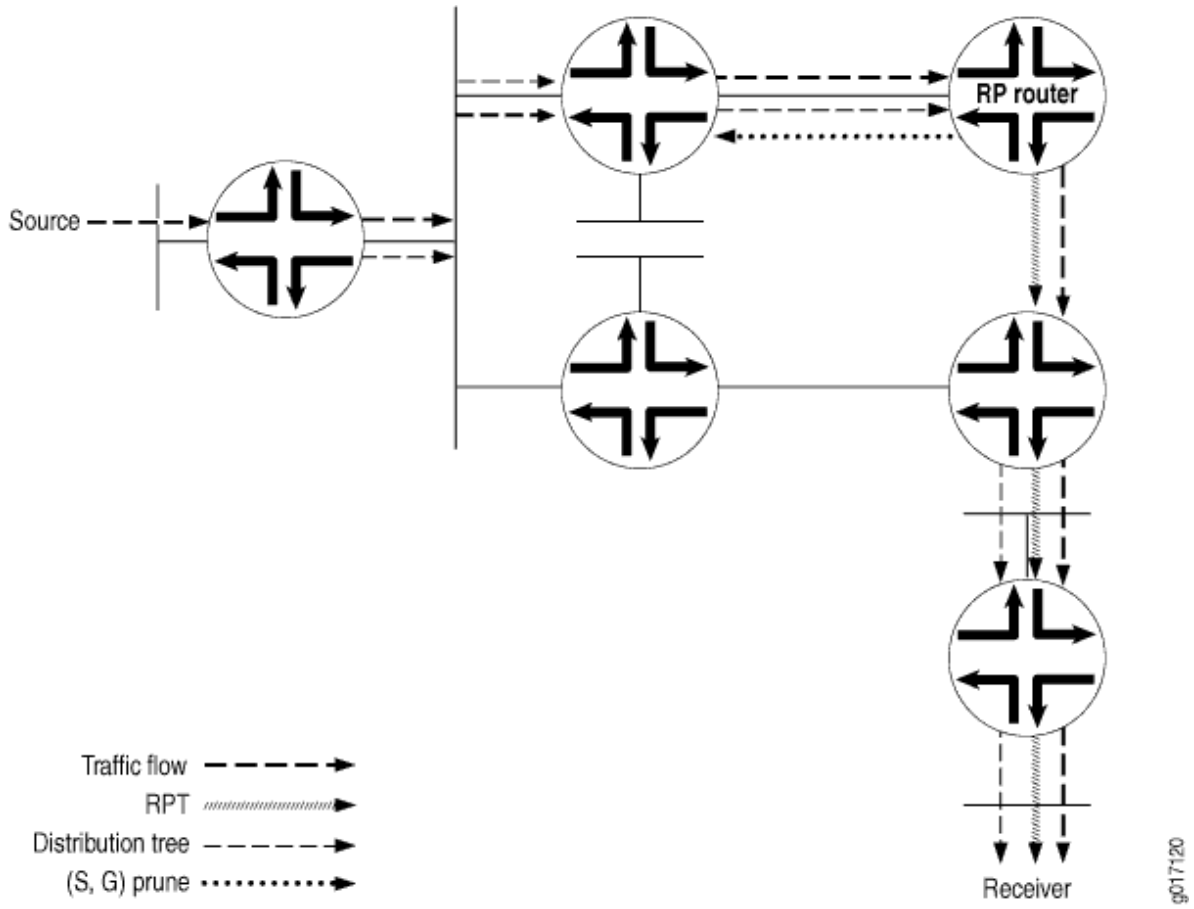
5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see [Figure 52 on page 407](#)).

Figure 52: RP Router Receives PIM Prune Message



- To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see [Figure 53 on page 408](#)).


Figure 53: RP Router Sends a PIM Prune Message to the Source DR



7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see [Figure 54 on page 409](#)).

g017120

Figure 54: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router



SPT Cutover Control

In some cases, the last-hop router needs to stay on the shared tree to the RP and not transition to a direct SPT to the source. You might not want the last-hop router to transition when, for example, a low-bandwidth multicast stream is forwarded from the RP to a last-hop router. All routers between last hop and source must maintain and refresh the SPT state. This can become a resource-intensive activity that does not add much to the network efficiency for a particular pair of source and multicast group addresses.

In these cases, you configure an SPT threshold policy on the last-hop router to control the transition to a direct SPT. An SPT cutover threshold of infinity applied to a source-group address pair means the last-hop router will never transition to a direct SPT. For all other source-group address pairs, the last-hop router transitions immediately to a direct SPT rooted at the source DR.

Example: Configuring the PIM Assert Timeout

IN THIS SECTION

- Requirements | 410
- Overview | 410
- Configuration | 412

This example shows how to configure the timeout period for a PIM assert forwarder.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure PIM Sparse Mode on the interfaces. See "[Enabling PIM Sparse Mode](#)" on page 316.

Overview

IN THIS SECTION

- Topology | 411

The role of PIM assert messages is to determine the forwarder on a network with multiple routers. The forwarder is the router that forwards multicast packets to a network with multicast group members. The forwarder is generally the same as the PIM DR.

A router sends an assert message when it receives a multicast packet on an interface that is listed in the outgoing interface list of the matching routing entry. Receiving a message on an outgoing interface is an indication that more than one router forwards the same multicast packets to a network.

In [Figure 55 on page 412](#), both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the

Ethernet network. An assert message contains, in addition to a source address and group address, a unicast cost metric for sending packets to the source, and a preference metric for the unicast cost. The preference metric expresses a preference between unicast routing protocols. The routing device with the smallest preference metric becomes the forwarder (also called the assert winner). If the preference metrics are equal, the device that sent the lowest unicast cost metric becomes the forwarder. If the unicast metrics are also equal, the routing device with the highest IP address becomes the forwarder. After the transmission of assert messages, only the forwarder continues to forward messages on the network.

When an assert message is received and the RPF neighbor is changed to the assert winner, the assert timer is set to an assert timeout period. The assert timeout period is restarted every time a subsequent assert message for the route entry is received on the incoming interface. When the assert timer expires, the routing device resets its RPF neighbor according to its unicast routing table. Then, if multiple forwarders still exist, the forwarders reenter the assert message cycle. In effect, the assert timeout period determines how often multicast routing devices enter a PIM assert message cycle.

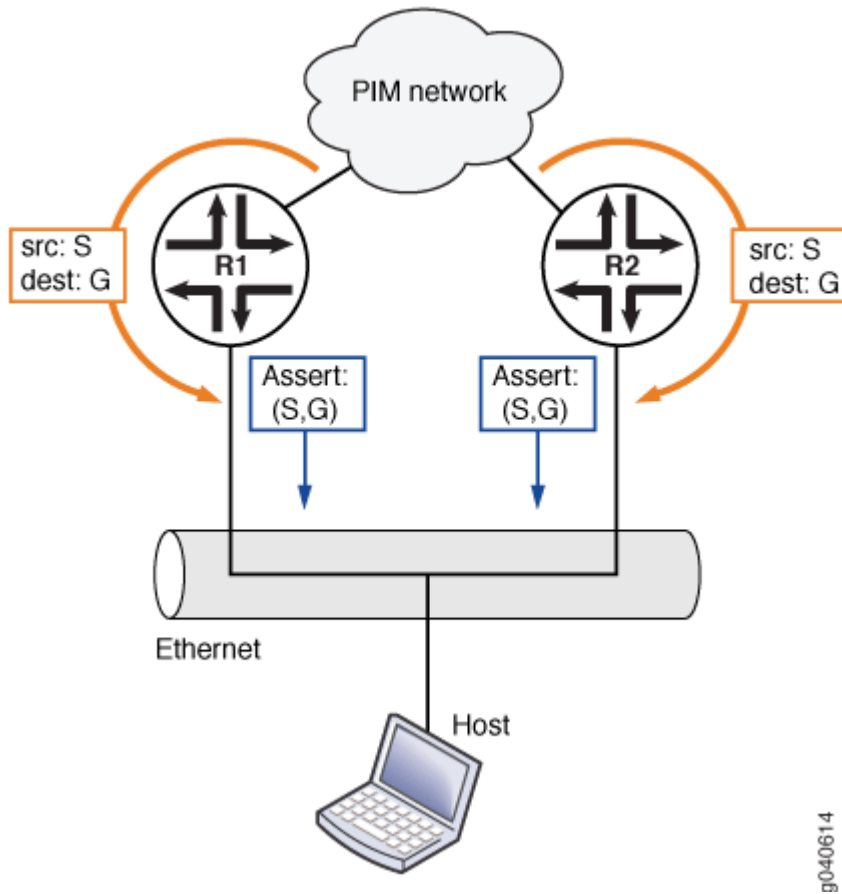
The range is from 5 through 210 seconds. The default is 180 seconds.

Assert messages are useful for LANs that connect multiple routing devices and no hosts.

Topology

[Figure 55 on page 412](#) shows the topology for this example.

Figure 55: PIM Assert Topology



g040614

Configuration

IN THIS SECTION

- Procedure | 412

Procedure

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 [Junos OS CLI User Guide](#).

To configure an assert timeout:

1. Configure the timeout period, in seconds.

```
[edit protocols pim]
user@host# set assert-timeout 60
```

2. (Optional) Trace assert messages.

```
[edit protocols pim]
user@host# set traceoptions file PIM.log
user@host# set traceoptions flag assert detail
```

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

```
user@host# commit
```

4. To verify the configuration, run the following commands:

- **show pim join**
- **show pim statistics**

SEE ALSO

[Configuring PIM Trace Options | 286](#)

[SPT Cutover | 404](#)

[SPT Cutover Control | 409](#)

Example: Configuring the PIM SPT Threshold Policy

IN THIS SECTION

- [Requirements | 414](#)
- [Overview | 414](#)
- [Configuration | 416](#)
- [Verification | 418](#)

This example shows how to apply a policy that suppresses the transition from the rendezvous-point tree (RPT) rooted at the RP to the shortest-path tree (SPT) rooted at the source.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure PIM Sparse Mode on the interfaces. See "[Enabling PIM Sparse Mode](#)" on page 316.

Overview

IN THIS SECTION

- [Topology | 415](#)

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or through an SPT rooted at the source. In some cases, the last-hop routing device needs to stay on the shared RPT to the RP and not transition to a direct SPT to the source. Receiving the multicast data traffic on SPT is optimal but introduces more state in the network, which might not be desirable in some multicast deployments. Ideally, low-bandwidth multicast streams can be forwarded on the RPT, and high-bandwidth streams can use the SPT. This example shows how to configure such a policy.

This example includes the following settings:

- **spt-threshold**—Enables you to configure an SPT threshold policy on the last-hop routing device to control the transition to a direct SPT. When you include this statement in the main PIM instance, the PE router stays on the RPT for control traffic.
- **infinity**—Applies an SPT cutover threshold of infinity to a source-group address pair, so that the last-hop routing device never transitions to a direct SPT. For all other source-group address pairs, the last-hop routing device transitions immediately to a direct SPT rooted at the source DR. This statement must reference a properly configured policy to set the SPT cutover threshold for a particular source-group pair to infinity. The use of values other than infinity for the SPT threshold is not supported. You can configure more than one policy.

- **policy-statement**—Configures the policy. The simplest type of SPT threshold policy uses a route filter and source address filter to specify the multicast group and source addresses and to set the SPT threshold for that pair of addresses to infinity. The policy is applied to the main PIM instance.

This example sets the SPT transition value for the source-group pair 10.10.10.1 and 224.1.1.1 to infinity. When the policy is applied to the last-hop router, multicast traffic from this source-group pair never transitions to a direct SPT to the source. Traffic will continue to arrive through the RP. However, traffic for any other source-group address combination at this router transitions to a direct SPT to the source.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.
- When the policy is configured for the first time, the routing device continues to transition to the direct SPT for the source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- If you do not clear the PIM-join state when you apply the infinity policy configuration for the first time, you must apply it before the PE router is brought up.
- When the policy is deleted for a source-group address pair for the first time, the routing device does not transition to the direct SPT for that source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- When the policy is changed for a source-group address pair for the first time, the routing device does not use the new policy until the PIM-join state is cleared with the **clear pim join** command.

Topology

Configuration

IN THIS SECTION

- Procedure | 416
- Results | 417

Procedure

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.

```
[edit]
set policy-options policy-statement spt-infinity-policy term one from route-filter 224.1.1.1/32
exact
set policy-options policy-statement spt-infinity-policy term one from source-address-filter
10.10.10.1/32 exact
set policy-options policy-statement spt-infinity-policy term one then accept
set policy-options policy-statement spt-infinity-policy term two then reject
set protocols pim spt-threshold infinity spt-infinity-policy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the [Junos OS CLI User Guide](#).

To configure an SPT threshold policy:

1. Apply the policy.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set spt-threshold infinity spt-infinity-policy
```

```
[edit protocols pim]
user@host# exit
```

2. Configure the policy.

```
[edit]
user@host# edit policy-options policy-statement spt-infinity-policy
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from route-filter 224.1.1.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from source-address-filter 10.10.10.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one then accept
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term two then reject
[edit policy-options policy-statement spt-infinity-policy]
user@host# exit
policy-statement {
```

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

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

4. Clear the PIM join cache to force the configuration to take effect.

```
[edit]
user@host# run clear pim join
```

Results

Confirm your configuration by entering the **show policy-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement spt-infinity-policy {
  term one {
```

```
    from {
        route-filter 224.1.1.1/32 exact;
        source-address-filter 10.10.10.1/32 exact;
    }
    then accept;
}
term two {
    then reject;
}
}
```

```
user@host# show protocols
pim {
    spt-threshold {
        infinity spt-infinity-policy;
    }
}
```

Verification

To verify the configuration, run the **show pim join** command.

SEE ALSO

[SPT Cutover Control | 409](#)

RELATED DOCUMENTATION

[Configuring PIM Auto-RP](#)

[Configuring PIM Bootstrap Router | 365](#)

[Configuring PIM Dense Mode | 299](#)

[Configuring a Designated Router for PIM | 425](#)

[Configuring PIM Filtering | 376](#)

[Example: Configuring Nonstop Active Routing for PIM | 515](#)

[Configuring PIM Sparse-Dense Mode | 304](#)

[Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol | 497](#)

[Configuring Basic PIM Settings](#)

Disabling PIM

IN THIS SECTION

- [Disabling the PIM Protocol | 419](#)
- [Disabling PIM on an Interface | 420](#)
- [Disabling PIM for a Family | 421](#)
- [Disabling PIM for a Rendezvous Point | 422](#)

By default, when you enable the PIM protocol it applies to the specified interface only. To enable PIM for all interfaces, include the `all` parameter (for example, set `protocol pim interface all`). You can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the `[edit protocols pim interface interface-name family]` hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the `[edit protocols pim interface interface-name family]` hierarchy level, but is disabled at the `[edit protocols pim interface interface-name]` hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the `[edit protocols pim interface interface-name family]` hierarchy level or the `[edit protocols pim interface interface-name]` hierarchy level, but is disabled at the `[edit protocols pim]` hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

Disabling the PIM Protocol

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the `[edit protocols pim]` hierarchy level:

```
[edit protocols]
pim {
```

```

    disable;
}

```

To disable the PIM protocol:

1. Include the `disable` statement.

```

user@host# set protocols pim disable

```

2. (Optional) Verify your configuration settings before committing them by using the `show protocols pim` command.

```

user@host# run show protocols pim

```

SEE ALSO

disable (PIM)

pim

Disabling PIM on an Interface

You can disable the PIM protocol on a per-interface basis. This is accomplished at the `[edit protocols pim interface interface-name]` hierarchy level:

```

[edit protocols]
pim {
    interface interface-name {
        disable;
    }
}

```

To disable PIM on an interface:

1. Include the `disable` statement.

```

user@host# set protocols pim interface fe-0/1/0 disable

```

2. (Optional) Verify your configuration settings before committing them by using the `show protocols pim` command.

```
user@host# run show protocols pim
```

SEE ALSO

disable (PIM)

pim

Disabling PIM for a Family

You can disable the PIM protocol on a per-family basis. This is accomplished at the `[edit protocols pim family]` hierarchy level:

```
[edit protocols]
pim {
  family inet {
    disable;
  }
  family inet6 {
    disable;
  }
}
```

To disable PIM for a family:

1. Include the `disable` statement.

```
user@host# set protocols pim family inet disable
user@host# set protocols pim family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the `show protocols pim` command.

```
user@host# run show protocols pim
```

SEE ALSO[disable \(PIM\)](#)[family \(Protocols PIM\)](#)[pim](#)**Disabling PIM for a Rendezvous Point**

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the [edit protocols pim rp local family] hierarchy level:

```
[edit protocols]
pim {
  rp {
    local {
      family inet {
        disable;
      }
      family inet6 {
        disable;
      }
    }
  }
}
```

To disable PIM for an RP family:

1. Use the disable statement.

```
user@host# set protocols pim rp local family inet disable
user@host# set protocols pim rp local family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the show protocols pim command.

```
user@host# run show protocols pim
```

SEE ALSO[family \(Local RP\)](#)

| *pim*

Configuring Designated Routers

IN THIS CHAPTER

- [Understanding Designated Routers | 424](#)
- [Configuring a Designated Router for PIM | 425](#)

Understanding Designated Routers

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers (DRs) to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.



NOTE: DR priority is specific to PIM sparse mode; as per RFC 3973, PIM DR priority cannot be configured explicitly in PIM Dense Mode (PIM-DM) in IGMPv2 – PIM-DM only support DRs with IGMPv1.

Configuring a Designated Router for PIM

IN THIS SECTION

- [Configuring Interface Priority for PIM Designated Router Selection | 425](#)
- [Configuring PIM Designated Router Election on Point-to-Point Links | 427](#)

Configuring Interface Priority for PIM Designated Router Selection

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) router learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has an equal probability (priority 1) of being selected as the DR, but you can change the value to increase or decrease the chances of a given DR being elected. A higher value corresponds to a higher priority, that is, greater chance of being elected. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.



NOTE: DR priority is specific to PIM sparse mode; as per RFC 3973, PIM DR priority cannot be configured explicitly in PIM Dense Mode (PIM-DM) in IGMPv2 – PIM-DM only support DRs with IGMPv1.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface ge-0/0/0.0]
user@host# set priority 5
```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the `show pim neighbors detail` command.

```
user@host> show pim neighbors detail
```

```
Instance: PIM.master
Interface: ge-0/0/0.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 5
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

SEE ALSO

[Configuring PIM Designated Router Election on Point-to-Point Links](#)

[Understanding PIM Sparse Mode | 307](#)

show pim neighbors

[Configuring Basic PIM Settings](#)

[Configuring PIM Sparse-Dense Mode | 304](#)

[Configuring PIM Dense Mode | 299](#)

[Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol | 497](#)

[Configuring PIM Auto-RP](#)

[Configuring PIM Filtering | 376](#)

priority (PIM Interfaces)

Configuring PIM Designated Router Election on Point-to-Point Links

In PIM sparse mode, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routers might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages. (DR priority is specific to PIM sparse mode; as per RFC 3973, PIM DR priority cannot be configured explicitly in PIM Dense Mode (PIM-DM) in IGMPv2 – PIM-DM only support DRs with IGMPv1.)

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routers, configure the router globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set dr-election-on-p2p
```

2. Verify the configuration by checking the **State** field in the output of the `show pim interfaces` command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.
3. If the `show pim interfaces` command continues to report the P2P state, consider running the `restart routing` command on both routers on the point-to-point link. Then recheck the state.



CAUTION: Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

```
[edit]
user@host# run restart routing
```

SEE ALSO

[Understanding PIM Sparse Mode | 307](#)

Configuring Interface Priority for PIM Designated Router Selection

show pim interfaces

Receiving Content Directly from the Source with SSM

IN THIS CHAPTER

- [Understanding PIM Source-Specific Mode | 429](#)
- [Example: Configuring Source-Specific Multicast | 434](#)
- [Example: Configuring PIM SSM on a Network | 451](#)
- [Example: Configuring an SSM-Only Domain | 454](#)
- [Example: Configuring SSM Mapping | 455](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 458](#)
- [Example: Configuring SSM Maps for Different Groups to Different Sources | 464](#)

Understanding PIM Source-Specific Mode

IN THIS SECTION

- [Any Source Multicast \(ASM\) was the Original Multicast | 430](#)
- [Source Discovery in Sparse Mode vs Dense Mode | 430](#)
- [PIM SSM is a Subset of PIM Sparse Mode | 430](#)
- [Why Use PIM SSM | 430](#)
- [PIM Terminology | 431](#)
- [How PIM SSM Works | 431](#)
- [Using PIM SSM | 433](#)

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode

functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

Any Source Multicast (ASM) was the Original Multicast

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of **source discovery** is a required function of the network itself.

Source Discovery in Sparse Mode vs Dense Mode

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

PIM SSM is a Subset of PIM Sparse Mode

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to **subscribers** (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

Why Use PIM SSM

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to **subscribers** (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM Terminology

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in [Table 14 on page 431](#).

Table 14: ASM and SSM Terminology

Term	Any-Source Multicast	Source-Specific Multicast
Address identifier	G	S,G
Address designation	group	channel
Receiver operations	join, leave	subscribe, unsubscribe
Group address range	224/4 excluding 232/8	224/4 (guaranteed only for 232/8)

Although PIM SSM describes receiver operations as **subscribe** and **subscriber**, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.

How PIM SSM Works


PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode

functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the [edit routing-options multicast] hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see [Figure 56 on page 432](#)). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in [Figure 56 on page 432](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 56: Receiver Announces Desire to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 57 on page 432](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 57: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 58 on page 433](#)).

Figure 58: (S,G) State Is Built Between the Source and the Receiver



Using PIM SSM

You can configure Junos OS to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.

To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.

RELATED DOCUMENTATION

[Source-Specific Multicast Groups Overview | 438](#)

[Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 458](#)

Example: Configuring Source-Specific Multicast

IN THIS SECTION

- [Understanding PIM Source-Specific Mode | 434](#)
- [Source-Specific Multicast Groups Overview | 438](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 439](#)
- [Example: Configuring an SSM-Only Domain | 445](#)
- [Example: Configuring PIM SSM on a Network | 445](#)
- [Example: Configuring SSM Mapping | 448](#)

Understanding PIM Source-Specific Mode

IN THIS SECTION

- [Any Source Multicast \(ASM\) was the Original Multicast | 434](#)
- [Source Discovery in Sparse Mode vs Dense Mode | 435](#)
- [PIM SSM is a Subset of PIM Sparse Mode | 435](#)
- [Why Use PIM SSM | 435](#)
- [PIM Terminology | 436](#)
- [How PIM SSM Works | 436](#)
- [Using PIM SSM | 438](#)

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

Any Source Multicast (ASM) was the Original Multicast

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the

locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of **source discovery** is a required function of the network itself.

Source Discovery in Sparse Mode vs Dense Mode

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

PIM SSM is a Subset of PIM Sparse Mode

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to **subscribers** (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

Why Use PIM SSM

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to **subscribers** (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM

bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM Terminology

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in [Table 15 on page 436](#).

Table 15: ASM and SSM Terminology

Term	Any-Source Multicast	Source-Specific Multicast
Address identifier	G	S,G
Address designation	group	channel
Receiver operations	join, leave	subscribe, unsubscribe
Group address range	224/4 excluding 232/8	224/4 (guaranteed only for 232/8)

Although PIM SSM describes receiver operations as **subscribe** and **subscriber**, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.


How PIM SSM Works

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see [Figure 59 on page 437](#)). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in [Figure 59 on page 437](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 59: Receiver Announces Desire to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 60 on page 437](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 60: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 61 on page 438](#)).

Figure 61: (S,G) State Is Built Between the Source and the Receiver



Using PIM SSM

You can configure Junos OS to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.

To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.

SEE ALSO

[Source-Specific Multicast Groups Overview | 438](#)

[Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 458](#)

Source-Specific Multicast Groups Overview

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM implemented in Junos OS has the efficient explicit join procedures of Protocol Independent Multicast (PIM) sparse mode but eliminates the immediate shared tree and rendezvous point (RP) procedures using (*,G) pairs. The (*) is a wildcard referring to any source sending to group G, and "G" refers to the IP multicast group. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The "S" refers to the source's unicast IP address, and the "G" refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. Although ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click a link in a browser, the receiver

is notified about the group information, but not the source information. With SSM, the client receives both source and group information.

SSM is ideal for one-to-many multicast services such as network entertainment channels. However, many-to-many multicast services might require ASM.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3.

SSM mapping allows operators to support an SSM network without requiring all hosts to support IGMPv3. This support exists in static (S,G) configurations, but SSM mapping also supports dynamic per-source group state information, which changes as hosts join and leave the group using IGMP.

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through the Multicast Source Discovery Protocol (MSDP).

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

IN THIS SECTION

- [Requirements | 439](#)
- [Overview | 440](#)
- [Configuration | 442](#)
- [Verification | 444](#)

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

Requirements

Before you begin, configure the router interfaces.

Overview


IN THIS SECTION

- Topology | 441

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see [Figure 62 on page 440](#)). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in [Figure 62 on page 440](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 62: Receiver Sends Messages to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 63 on page 441](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 63: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 64 on page 441).

Figure 64: (S,G) State Is Built Between the Source and the Receiver




SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

Topology

This example works with the simple RPF topology shown in Figure 65 on page 441.

Figure 65: Simple RPF Topology



Configuration

IN THIS SECTION

- Procedure | 442
- Results | 443

Procedure

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.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

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 [Junos OS CLI User Guide](#).

To configure an RPF policy:

1. Configure OSPF.

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

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols` and `show routing-options` commands.

```
user@host# show protocols
ospf {
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
```

```
    }
    interface all;
  }
}
pim {
  rp {
    local {
      address 10.255.72.46;
      group-ranges {
        239.0.0.0/24;
      }
    }
  }
  interface fe-1/0/0.0 {
    mode sparse;
  }
  interface lo0.0 {
    mode sparse;
  }
}
```

```
user@host# show routing-options
multicast {
  ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ];
  asm-override-ssm;
}
```

Verification

To verify the configuration, run the following commands:

- **show igmp group**
- **show igmp statistics**
- **show pim join**

SEE ALSO

| [Source-Specific Multicast Groups Overview](#) | 438

Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.


In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
  igmp {
    interface fe-0/1/2 {
      version 3;
    }
  }
}
```

Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in [Figure 66 on page 446](#).

Figure 66: Network on Which to Configure PIM SSM



This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

```
user@router1# set protocols igmp interface all version 3
user@router1# set protocols igmp interface fxp0.0 disable
```



NOTE: When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```
user@router1> show configuration protocol igmp
```

```
[edit protocols igmp]
interface all {
  version 3;
}
interface fxp0.0 {
  disable;
}
```

3. Use the `show igmp interface` command to verify that IGMP interfaces are configured.

```
user@router1> show igmp interface
Interface      State   Querier  Timeout Version Groups
```



```

fe-0/0/0.0    Up      198.51.100.245    213    3    0
fe-0/0/1.0    Up      198.51.100.241    220    3    0
fe-0/0/2.0    Up      198.51.100.237    218    3    0

```

Configured Parameters:

```

IGMP Query Interval (1/10 secs): 1250
IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550

```

- Use the `show pim join extensive` command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```

user@router2> show pim join extensive
232.1.1.1      10.4.1.2          sparse
  Upstream interface: fe-1/1/3.0
  Upstream State: Local Source
  Keepalive timeout: 209
  Downstream Neighbors:
    Interface: so-1/0/2.0
      10.10.71.1      State: Join  Flags: S   Timeout: 209

```

- Use the `show pim join extensive` command to verify the PIM join state on Router 1 (the router connected to the receiver).

```

user@router1> show pim join extensive
232.1.1.1      10.4.1.2          sparse
  Upstream interface: so-1/0/2.0
  Upstream State: Join to Source
  Keepalive timeout: 209
  Downstream Neighbors:
    Interface: fe-0/2/3.0
      10.3.1.1        State: Join  Flags: S   Timeout: Infinity

```



NOTE: IP version 6 (IPv6) multicast routers use the Multicast Listener Discovery (MLD) Protocol to manage the membership of hosts and routers in multicast groups and to

learn which groups have interested listeners for each attached physical networks. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol. MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

SEE ALSO

| [Example: Configuring SSM Mapping | 455](#)

Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

```
user@router1# set policy-options policy-statement ssm-policy-example term A from route-filter
232.1.1.1/32 exact
```

```

user@router1# set policy-options policy-statement ssm-policy-example term A then accept
user@router1# set policy-options policy-statement ssm-policy-example term B from route-filter
ff35::1/128 exact
user@router1# set policy-options policy-statement ssm-policy-example term B then accept

```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```

user@host> show configuration policy-options

```

```

[edit policy-options]
policy-statement ssm-policy-example {
  term A {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then accept;
  }
  term B {
    from {
      route-filter ff35::1/128 exact;
    }
    then accept;
  }
  then reject;
}

```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```

user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy ssm-policy-
example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy ssm-policy-
example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source 10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source 192.168.43.66

```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options
```

```
[edit routing-options]
multicast {
  ssm-map ssm-map-ipv6-example {
    policy ssm-policy-example;
    source [ fec0::1 fec0::12 ];
  }
  ssm-map ssm-map-ipv4-example {
    policy ssm-policy-example;
    source [ 10.10.10.4 192.168.43.66 ];
  }
}
```

We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol
```

```
[edit protocols]
igmp {
  interface fe-0/1/0.0 {
    ssm-map ssm-map-ipv4-example;
  }
}
mld {
  interface fe-0/1/1.0 {
    ssm-map ssm-map-ipv6-example;
```

```
}  
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```
user@host> show igmp interface fe-0/1/0.0  
Interface: fe-0/1/0.0  
  Querier: 192.168.224.28  
  State:          Up Timeout:  None Version:  2 Groups:  2  
  SSM Map: ssm-map-ipv4-example
```

```
user@host> show mld interface fe-0/1/1.0  
Interface: fe-0/1/1.0  
  Querier: fec0:0:0:0:1::12  
  State:          Up Timeout:  None Version:  2 Groups:  2  
  SSM Map: ssm-map-ipv6-example
```


RELATED DOCUMENTATION

| [Configuring Basic PIM Settings](#)

Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in [Figure 67 on page 452](#).

Figure 67: Network on Which to Configure PIM SSM



This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

```
user@router1# set protocols igmp interface all version 3
user@router1# set protocols igmp interface fxp0.0 disable
```



NOTE: When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the `show configuration protocol igmp` command to verify the IGMP protocol configuration.

```
user@router1> show configuration protocol igmp
```

```
[edit protocols igmp]
interface all {
  version 3;
}
interface fxp0.0 {
  disable;
}
```

3. Use the `show igmp interface` command to verify that IGMP interfaces are configured.

```
user@router1> show igmp interface
Interface      State   Querier  Timeout Version Groups
```

```

fe-0/0/0.0    Up      198.51.100.245    213    3    0
fe-0/0/1.0    Up      198.51.100.241    220    3    0
fe-0/0/2.0    Up      198.51.100.237    218    3    0

```

Configured Parameters:

```

IGMP Query Interval (1/10 secs): 1250
IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550

```

4. Use the `show pim join extensive` command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```

user@router2> show pim join extensive
232.1.1.1      10.4.1.2          sparse
  Upstream interface: fe-1/1/3.0
  Upstream State: Local Source
  Keepalive timeout: 209
  Downstream Neighbors:
    Interface: so-1/0/2.0
      10.10.71.1      State: Join  Flags: S   Timeout: 209

```

5. Use the `show pim join extensive` command to verify the PIM join state on Router 1 (the router connected to the receiver).

```

user@router1> show pim join extensive
232.1.1.1      10.4.1.2          sparse
  Upstream interface: so-1/0/2.0
  Upstream State: Join to Source
  Keepalive timeout: 209
  Downstream Neighbors:
    Interface: fe-0/2/3.0
      10.3.1.1        State: Join  Flags: S   Timeout: Infinity

```



NOTE: IP version 6 (IPv6) multicast routers use the Multicast Listener Discovery (MLD) Protocol to manage the membership of hosts and routers in multicast groups and to

learn which groups have interested listeners for each attached physical networks. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol. MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

RELATED DOCUMENTATION

| [Example: Configuring SSM Mapping | 455](#)

Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```



```

igmp {
  interface fe-0/1/2 {
    version 3;
  }
}

```

Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

```

user@router1# set policy-options policy-statement ssm-policy-example term A from route-filter
232.1.1.1/32 exact
user@router1# set policy-options policy-statement ssm-policy-example term A then accept
user@router1# set policy-options policy-statement ssm-policy-example term B from route-filter
ff35::1/128 exact
user@router1# set policy-options policy-statement ssm-policy-example term B then accept

```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```
user@host> show configuration policy-options
```

```
[edit policy-options]
policy-statement ssm-policy-example {
  term A {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then accept;
  }
  term B {
    from {
      route-filter ff35::1/128 exact;
    }
    then accept;
  }
  then reject;
}
```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy ssm-policy-
example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy ssm-policy-
example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source 10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source 192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options
```

```
[edit routing-options]
multicast {
  ssm-map ssm-map-ipv6-example {
    policy ssm-policy-example;
    source [ fec0::1 fec0::12 ];
  }
  ssm-map ssm-map-ipv4-example {
    policy ssm-policy-example;
    source [ 10.10.10.4 192.168.43.66 ];
  }
}
```

We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol
```

```
[edit protocols]
igmp {
  interface fe-0/1/0.0 {
    ssm-map ssm-map-ipv4-example;
  }
}
mld {
  interface fe-0/1/1.0 {
    ssm-map ssm-map-ipv6-example;
```

```
}
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```
user@host> show igmp interface fe-0/1/0.0
Interface: fe-0/1/0.0
  Querier: 192.168.224.28
  State:          Up Timeout:  None Version:  2 Groups:  2
  SSM Map: ssm-map-ipv4-example
```

```
user@host> show mld interface fe-0/1/1.0
Interface: fe-0/1/1.0
  Querier: fec0:0:0:0:1::12
  State:          Up Timeout:  None Version:  2 Groups:  2
  SSM Map: ssm-map-ipv6-example
```

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

IN THIS SECTION

- [Requirements | 459](#)
- [Overview | 459](#)
- [Configuration | 461](#)
- [Verification | 463](#)

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

Requirements

Before you begin, configure the router interfaces.

Overview


IN THIS SECTION

- Topology | 460

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see [Figure 68 on page 459](#)). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in [Figure 68 on page 459](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 68: Receiver Sends Messages to Join Group G and Source S




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 69 on page 460](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 69: Router 3 (Last-Hop Router) Joins the Source Tree



Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 70 on page 460).

Figure 70: (S,G) State Is Built Between the Source and the Receiver




SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

Topology

This example works with the simple RPF topology shown in Figure 71 on page 460.

Figure 71: Simple RPF Topology



Configuration

IN THIS SECTION

- Procedure | 461
- Results | 462

Procedure

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.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

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 [Junos OS CLI User Guide](#).

To configure an RPF policy:

1. Configure OSPF.

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

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show protocols` and `show routing-options` commands.

```
user@host# show protocols
ospf {
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
```



```
    }
    interface all;
  }
}
pim {
  rp {
    local {
      address 10.255.72.46;
      group-ranges {
        239.0.0.0/24;
      }
    }
  }
  interface fe-1/0/0.0 {
    mode sparse;
  }
  interface lo0.0 {
    mode sparse;
  }
}
```

```
user@host# show routing-options
multicast {
  ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ];
  asm-override-ssm;
}
```

Verification

To verify the configuration, run the following commands:

- **show igmp group**
- **show igmp statistics**
- **show pim join**

RELATED DOCUMENTATION

| [Source-Specific Multicast Groups Overview](#) | 438

Example: Configuring SSM Maps for Different Groups to Different Sources

IN THIS SECTION

- [Multiple SSM Maps and Groups for Interfaces | 464](#)
- [Example: Configuring Multiple SSM Maps Per Interface | 464](#)

Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

IN THIS SECTION

- [Requirements | 464](#)
- [Overview | 464](#)
- [Configuration | 465](#)
- [Verification | 468](#)

This example shows how to assign more than one SSM map to an IGMP interface.

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, `POLICY-ipv4-example1`, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 1	232.1.1.1	10.10.10.4, 192.168.43.66
POLICY-ipv4-example1 term 2	232.1.1.2	10.10.10.5, 192.168.43.67

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 465](#)
- [Procedure | 466](#)

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the [Junos OS CLI User Guide](#).

To configure this example, perform the following task:

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands 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.

```
set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter 232.1.1.1/32
exact
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source 10.10.10.4
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source 192.168.43.66
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter 232.1.1.2/32
exact
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source 10.10.10.5
```

```

set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source 192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1

```

Procedure

Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```

[edit policy-options policy-statement POLICY-ipv4-example1 term 1]
user@host# set from route-filter 232.1.1.1/32 exact
user@host# set then ssm-source 10.10.10.4
user@host# set then ssm-source 192.168.43.66
user@host# set then accept

```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```

[edit policy-options policy-statement POLICY-ipv4-example1 term 2]
user@host# set from route-filter 232.1.1.2/32 exact
user@host# set then ssm-source 10.10.10.5
user@host# set then ssm-source 192.168.43.67
user@host# set then accept

```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

```

[edit protocols igmp interface fe-0/1/0.0]
user@host# set ssm-map-policy POLICY-ipv4-example1

```

Results

After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
user@host# show policy-options
policy-statement POLICY-ipv4-example1 {
  term 1 {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then {
      ssm-source [ 10.10.10.4 192.168.43.66 ];
      accept;
    }
  }
  term 2{
    from {
      route-filter 232.1.1.2/32 exact;
    }
    then {
      ssm-source [ 10.10.10.5 192.168.43.67 ];
      accept;
    }
  }
}

user@host# show protocols
igmp {
  interface fe-0/1/0.0 {
    ssm-map-policy POLICY-ipv4-example1;
  }
}
```

Verification

IN THIS SECTION

- [Displaying Information About IGMP-Enabled Interfaces | 468](#)
- [Displaying the PIM Groups | 469](#)
- [Displaying the Entries in the IP Multicast Forwarding Table | 469](#)

Confirm that the configuration is working properly.

Displaying Information About IGMP-Enabled Interfaces

Purpose

Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

Action

Use the **show igmp interface** operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```
user@host> show igmp interface
Interface: fe-0/1/0.0
  Querier: 10.111.30.1
  State:      Up Timeout:  None Version:  2 Groups:    2
  SSM Map Policy: POLICY-ipv4-example1;

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

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

The command output displays the name of the IGMP logical interface (fe-0/1/0.0), which is the address of the routing device that has been elected to send membership queries and group information.

Displaying the PIM Groups

Purpose

Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

Action

Use the **show pim join extensive 232.1.1.1** operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

Displaying the Entries in the IP Multicast Forwarding Table

Purpose

Verify that the IP multicast forwarding table displays the multicast route state.

Action

Use the **show multicast route extensive** operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

RELATED DOCUMENTATION

[Example: Configuring Source-Specific Multicast](#)

[Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs | 648](#)

Minimizing Routing State Information with Bidirectional PIM

IN THIS CHAPTER

- [Example: Configuring Bidirectional PIM | 470](#)

Example: Configuring Bidirectional PIM

IN THIS SECTION

- [Understanding Bidirectional PIM | 470](#)
- [Example: Configuring Bidirectional PIM | 478](#)

Understanding Bidirectional PIM

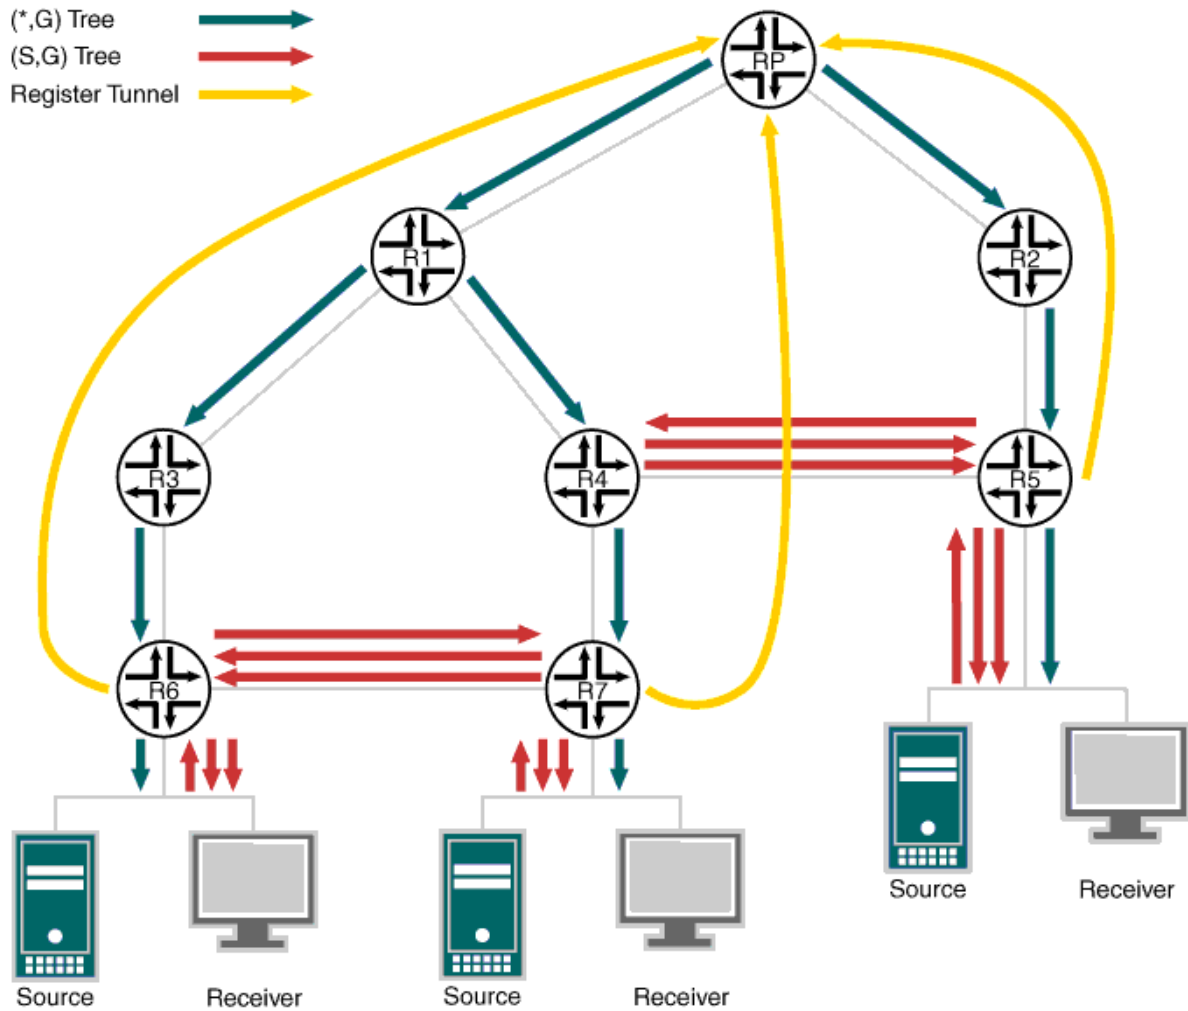
IN THIS SECTION

- [Designated Forwarder Election | 473](#)
- [Bidirectional PIM Modes | 474](#)
- [Bidirectional Rendezvous Points | 474](#)
- [PIM Bootstrap and Auto-RP Support | 475](#)
- [IGMP and MLD Support | 475](#)
- [Bidirectional PIM and Graceful Restart | 476](#)
- [Junos OS Enhancements to Bidirectional PIM | 476](#)
- [Limitations of Bidirectional PIM | 477](#)

Bidirectional PIM (PIM-Bidir) is specified by the IETF in RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*. It provides an alternative to other PIM modes, such as PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM source-specific multicast (SSM). In bidirectional PIM, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes the amount of PIM routing state information that must be maintained, which is especially important in networks with numerous and dispersed senders and receivers. For example, one important application for bidirectional PIM is distributed inventory polling. In many-to-many applications, a multicast query from one station generates multicast responses from many stations. For each multicast group, such an application generates a large number of (S,G) routes for each station in PIM-SM, PIM-DM, or SSM. The problem is even worse in applications that use bursty sources, resulting in frequently changing multicast tables and, therefore, performance problems in routers.

[Figure 72 on page 472](#) shows the traffic flows generated to deliver traffic for one group to and from three stations in a PIM-SM network.

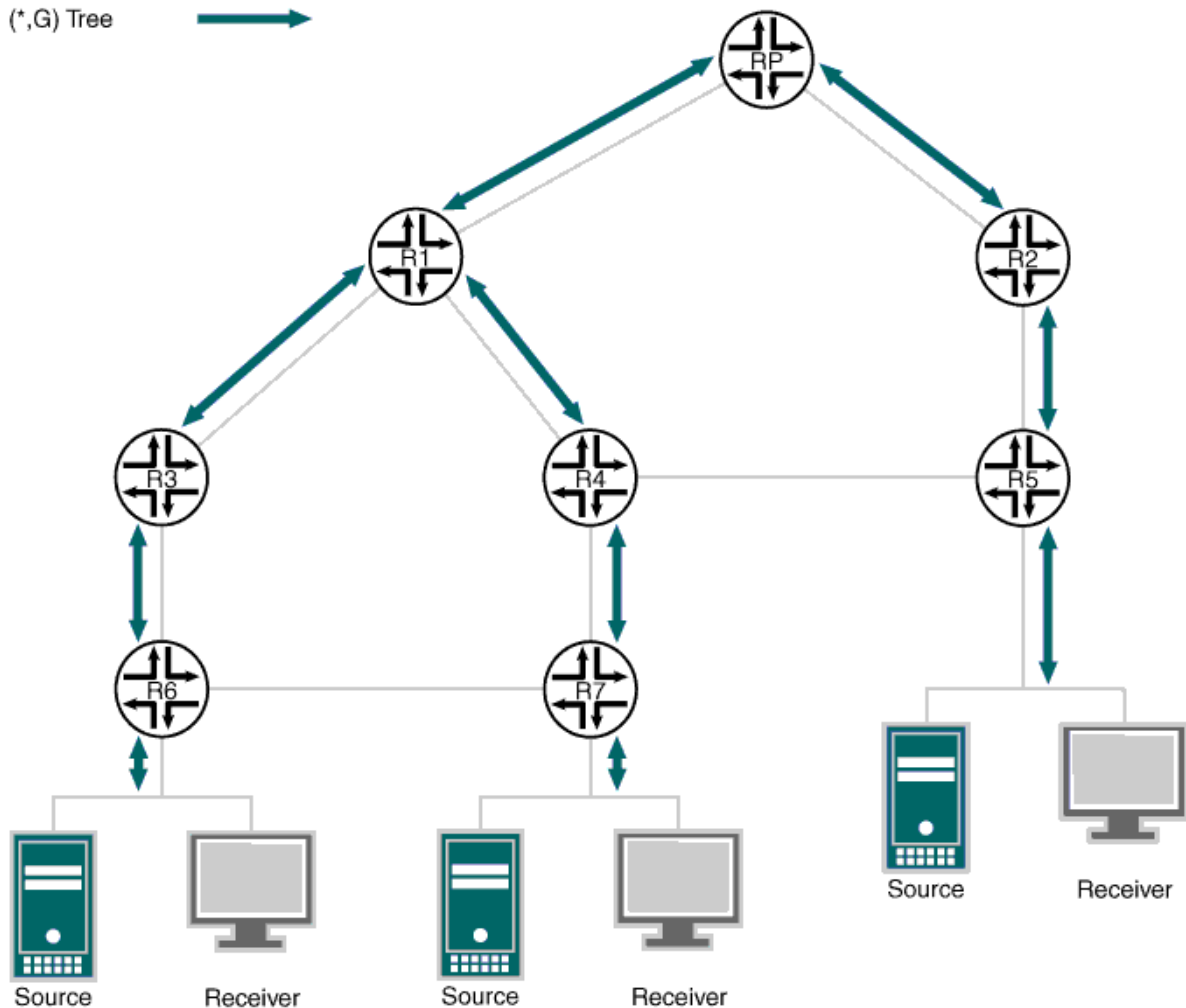
Figure 72: Example PIM Sparse-Mode Tree



Bidirectional PIM solves this problem by building only group-specific (*,G) state. Thus, only a single (*,G) route is needed for each group to deliver traffic to and from all the sources.

[Figure 73 on page 473](#) shows the traffic flows generated to deliver traffic for one group to and from three stations in a bidirectional PIM network.

Figure 73: Example Bidirectional PIM Tree



Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees (SPTs) as in PIM-SM and is therefore optimized for routing state size instead of path length. Bidirectional PIM routes are always wildcard-source (*,G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (*,G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM routes forward traffic from all sources and the RP. Thus, bidirectional PIM routers must have the ability to accept traffic on many potential incoming interfaces.

Designated Forwarder Election

To prevent forwarding loops, only one router on each link or subnet (including point-to-point links) is a designated forwarder (DF). The responsibilities of the DF are to forward downstream traffic onto the

link toward the receivers and to forward upstream traffic from the link toward the RP address. Bidirectional PIM relies on a process called DF election to choose the DF router for each interface and for each RP address. Each bidirectional PIM router in a subnet advertises its interior gateway protocol (IGP) unicast route to the RP address. The router with the best IGP unicast route to the RP address wins the DF election. Each router advertises its IGP route metrics in DF Offer, Winner, Backoff, and Pass messages.

Junos OS implements the DF election procedures as stated in RFC 5015, except that Junos OS checks RP unicast reachability before accepting incoming DF messages. DF messages for unreachable rendezvous points are ignored.

Bidirectional PIM Modes

In the Junos OS implementation, there are two modes for bidirectional PIM: bidirectional-sparse and bidirectional-sparse-dense. The differences between bidirectional-sparse and bidirectional-sparse-dense modes are the same as the differences between sparse mode and sparse-dense mode. Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Use bidirectional-sparse-dense mode when you have a mix of bidirectional groups, sparse groups, and dense groups in your network. One typical scenario for this is the use of auto-RP, which uses dense-mode flooding to bootstrap itself for sparse mode or bidirectional mode. In general, the dense groups could be for any flows that the network design requires to be flooded.

Each group-to-RP mapping is controlled by the RP `group-ranges` statement and the `ssm-groups` statement.

The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:

- bidirectional-sparse—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.
- bidirectional-sparse-dense—Use if multicast groups, except those that are specified in the `dense-groups` statement, are operating in bidirectional, sparse, or SSM mode.

Bidirectional Rendezvous Points

You can configure group-range-to-RP mappings network-wide statically, or only on routers connected to the RP addresses and advertise them dynamically. Unlike rendezvous points for PIM-SM, which must de-encapsulate PIM Register messages and perform other specific protocol actions, bidirectional PIM rendezvous points implement no specific functionality. RP addresses are simply locations in the network to rendezvous toward. In fact, RP addresses need not be loopback interface addresses or even be addresses configured on any router, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable router and advertised to the network.

Thus, for bidirectional PIM, there is no meaningful distinction between static and local RP addresses. Therefore, bidirectional PIM rendezvous points are configured at the `[edit protocols pim rp bidirectional]` hierarchy level, not under `static` or `local`.

The settings at the `[edit protocol pim rp bidirectional]` hierarchy level function like the settings at the `[edit protocols pim rp local]` hierarchy level, except that they create bidirectional PIM RP state instead of PIM-SM RP state.

Where only a single local RP can be configured, multiple bidirectional rendezvous points can be configured having group ranges that are the same, different, or overlapping. It is also permissible for a group range or RP address to be configured as bidirectional and either static or local for sparse-mode.

If a bidirectional PIM RP is configured without a group range, the default group range is 224/4 for IPv4. For IPv6, the default is ff00::/8. You can configure a bidirectional PIM RP group range to cover an SSM group range, but in that case the SSM or DM group range takes precedence over the bidirectional PIM RP configuration for those groups. In other words, because SSM always takes precedence, it is not permitted to have a bidirectional group range equal to or more specific than an SSM or DM group range.

PIM Bootstrap and Auto-RP Support

Group ranges for the specified RP address are flagged by PIM as bidirectional PIM group-to-RP mappings and, if configured, are advertised using PIM bootstrap or auto-RP. Dynamic advertisement of bidirectional PIM-flagged group-to-RP mappings using PIM bootstrap, and auto-RP is controlled as normal using the `bootstrap` and `auto-rp` statements.

Bidirectional PIM RP addresses configured at the `[edit protocols pim rp bidirectional address]` hierarchy level are advertised by auto-RP or PIM bootstrap if the following prerequisites are met:

- The routing instance must be configured to advertise candidate rendezvous points by way of auto-RP or PIM bootstrap, and an auto-RP mapping agent or bootstrap router, respectively, must be elected.
- The RP address must either be configured locally on an interface in the routing instance, or the RP address must belong to a subnet connected to an interface in the routing instance.

IGMP and MLD Support

Internet Group Management Protocol (IGMP) version 1, version 2, and version 3 are supported with bidirectional PIM. Multicast Listener Discovery (MLD) version 1 and version 2 are supported with bidirectional PIM. However, in all cases, only anysource multicast (ASM) state is supported for bidirectional PIM membership.

The following rules apply to bidirectional PIM:

- IGMP and MLD (*,G) membership reports trigger the PIM DF to originate bidirectional PIM (*,G) join messages.
- IGMP and MLD (S,G) membership reports do not trigger the PIM DF to originate bidirectional PIM (*,G) join messages.

Bidirectional PIM and Graceful Restart

Bidirectional PIM accepts packets for a bidirectional route on multiple interfaces. This means that some topologies might develop multicast routing loops if all PIM neighbors are not synchronized with regard to the identity of the designated forwarder (DF) on each link. If one router is forwarding without actively participating in DF elections, particularly after unicast routing changes, multicast routing loops might occur.

If graceful restart for PIM is enabled and bidirectional PIM is enabled, the default graceful restart behavior is to continue forwarding packets on bidirectional routes. If the gracefully restarting router was serving as a DF for some interfaces to rendezvous points, the restarting router sends a DF Winner message with a metric of 0 on each of these RP interfaces. This ensures that a neighbor router does not become the DF due to unicast topology changes that might occur during the graceful restart period. Sending a DF Winner message with a metric of 0 prevents another PIM neighbor from assuming the DF role until after graceful restart completes. When graceful restart completes, the gracefully restarted router sends another DF Winner message with the actual converged unicast metric.

The `no-bidirectional-mode` statement at the `[edit protocols pim graceful-restart]` hierarchy level overrides the default behavior and disables forwarding for bidirectional PIM routes during graceful restart recovery, both in cases of simple routing protocol process (rpd) restart and *graceful Routing Engine switchover*. This *configuration statement* provides a very conservative alternative to the default graceful restart behavior for bidirectional PIM routes. The reason to discontinue forwarding of packets on bidirectional routes is that the continuation of forwarding might lead to short-duration multicast loops in rare double-failure circumstances.

Junos OS Enhancements to Bidirectional PIM

In addition to the functionality specified in RFC 5015, the following functions are included in the Junos OS implementation of bidirectional PIM:

- Source-only branches without PIM join state
- Support for both IPv4 and IPv6 domain and multicast addresses
- Nonstop routing (NSR) for bidirectional PIM routes
- Support for bidirectional PIM in logical systems
- Support for non-forwarding and virtual router instances

The following caveats are applicable for the bidirectional PIM configuration on the PTX5000:

- PTX5000 routers can be configured both as a bidirectional PIM rendezvous point and the source node.

- For PTX5000 routers, you can configure the `auto-rp` statement at the `[edit protocols pim rp]` or the `[edit routing-instances routing-instance-name protocols pim rp]` hierarchy level with the `mapping` option, but not the `announce` option.

Limitations of Bidirectional PIM

The Junos OS implementation of bidirectional PIM does not support the following functionality:

Starting with Release 12.2, Junos OS extends the *nonstop active routing* PIM support to draft-rosen MVPNs.

PTX5000 routers do not support nonstop active routing or in-service software upgrade (ISSU) in Junos OS Release 13.3.

Nonstop active routing PIM support for draft-rosen MVPNs enables nonstop active routing-enabled devices to preserve draft-rosen MPVN-related information—such as default and data MDT states—across switchovers.

- SNMP for bidirectional PIM.
- Graceful Routing Engine switchover is configurable with bidirectional PIM enabled, but bidirectional routes do not forward packets during the switchover.
- Multicast VPNs (Draft Rosen and NextGen).

The bidirectional PIM protocol does not support the following functionality:

- Embedded RP
- Anycast RP

SEE ALSO

[Example: Configuring Bidirectional PIM | 470](#)

[Configuring PIM Auto-RP](#)

[Junos OS Multicast Protocols User Guide](#)

[Configuring PIM Bootstrap Properties for IPv4 or IPv6 | 367](#)

[Junos OS Multicast Protocols User Guide](#)

Example: Configuring Bidirectional PIM

IN THIS SECTION

- [Requirements | 478](#)
- [Overview | 478](#)
- [Configuration | 481](#)
- [Verification | 488](#)

This example shows how to configure bidirectional PIM, as specified in RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*.

Requirements

This example uses the following hardware and software components:

- Eight Juniper Networks routers that can be M120, M320, MX Series, or T Series platforms. To support bidirectional PIM, M Series platforms must have I-chip FPCs. M7i, M10i, M40e, and other older M Series routers do not support bidirectional PIM.
- Junos OS Release 12.1 or later running on all eight routers.

Overview

IN THIS SECTION

- [Topology Diagram | 479](#)

Compared to PIM sparse mode, bidirectional PIM requires less PIM router state information. Because less state information is required, bidirectional PIM scales well and is useful in deployments with many dispersed sources and receivers.

In this example, two rendezvous points are configured statically. One RP is configured as a phantom RP. A phantom RP is an RP address that is a valid address on a subnet, but is not assigned to a PIM router interface. The subnet must be reachable by the bidirectional PIM routers in the network. For the other (non-phantom) RP in this example, the RP address is assigned to a PIM router interface. It can be

assigned to either the loopback interface or any physical interface on the router. In this example, it is assigned to a physical interface.

OSPF is used as the interior gateway protocol (IGP) in this example. The OSPF metric determines the designated forwarder (DF) election process. In bidirectional PIM, the DF establishes a loop-free shortest-path tree that is rooted at the RP. On every network segment and point-to-point link, all PIM routers participate in DF election. The procedure selects one router as the DF for every RP of bidirectional groups. This router forwards multicast packets received on that network upstream to the RP. The DF election uses the same tie-break rules used by PIM assert processes.

This example uses the default DF election parameters. Optionally, at the **[edit protocols pim interface (interface-name | all) bidirectional]** hierarchy level, you can configure the following parameters related to the DF election:

- The robustness-count is the minimum number of DF election messages that must be lost for election to fail.
- The offer period is the interval to wait between repeated DF Offer and Winner messages.
- The backoff period is the period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

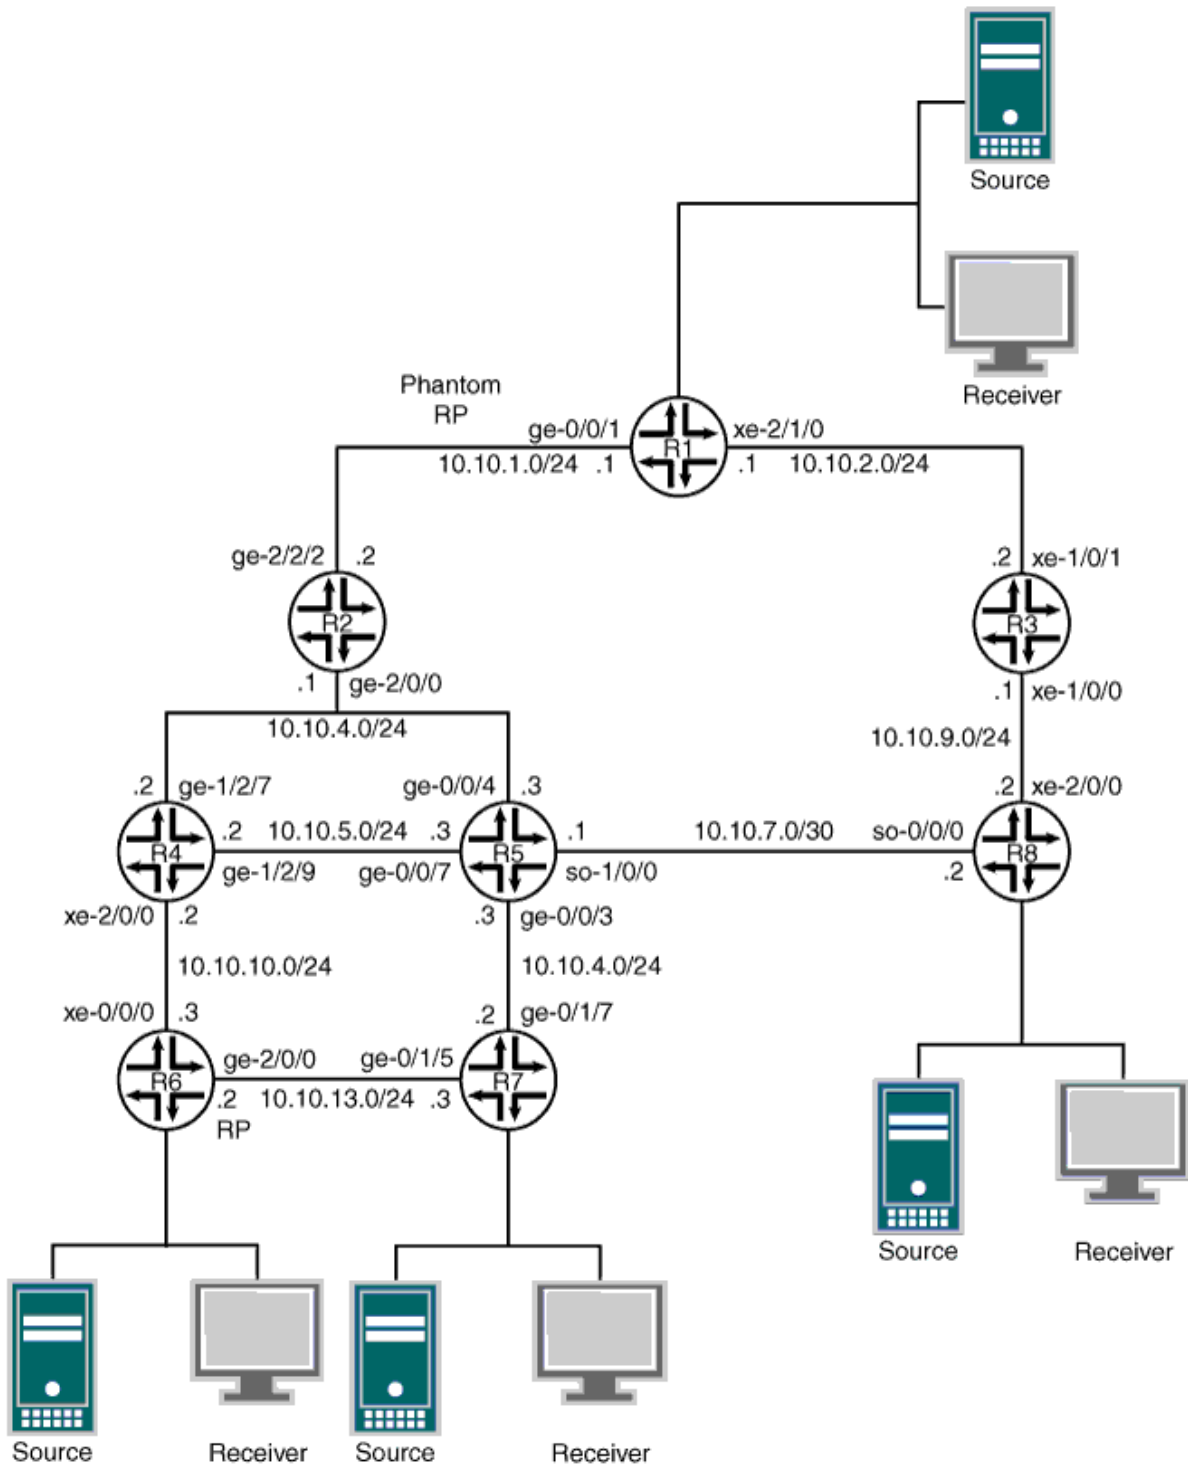
This example uses bidirectional-sparse-dense mode on the interfaces. The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:

- **bidirectional-sparse**—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.
- **bidirectional-sparse-dense**—Use if multicast groups, except those that are specified in the **dense-groups** statement, are operating in bidirectional, sparse, or SSM mode.

Topology Diagram

[Figure 74 on page 480](#) shows the topology used in this example.

Figure 74: Bidirectional PIM with Statically Configured Rendezvous Points



9680896

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 481](#)
- [Router R1 | 485](#)
- [Results | 486](#)

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 R1

```
set interfaces ge-0/0/1 unit 0 family inet address 10.10.1.1/24
set interfaces xe-2/1/0 unit 0 family inet address 10.10.2.1/24
set interfaces lo0 unit 0 family inet address 10.255.11.11/32
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface xe-2/1/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim interface ge-0/0/1.0 mode bidirectional-sparse-dense
set protocols pim interface xe-2/1/0.0 mode bidirectional-sparse-dense
```

Router R2

```
set interfaces ge-2/0/0 unit 0 family inet address 10.10.4.1/24
set interfaces ge-2/2/2 unit 0 family inet address 10.10.1.2/24
set interfaces lo0 unit 0 family inet address 10.255.22.22/32
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```

set protocols ospf area 0.0.0.0 interface ge-2/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim interface fxp0.0 disable
set protocols pim interface ge-2/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-2/2/2.0 mode bidirectional-sparse-dense

```

Router R3

```

set interfaces xe-1/0/0 unit 0 family inet address 10.10.9.1/24
set interfaces xe-1/0/1 unit 0 family inet address 10.10.2.2/24
set interfaces lo0 unit 0 family inet address 10.255.33.33/32
set protocols ospf area 0.0.0.0 interface xe-1/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface xe-1/0/0.0
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim interface xe-1/0/1.0 mode bidirectional-sparse-dense
set protocols pim interface xe-1/0/0.0 mode bidirectional-sparse-dense

```

Router R4

```

set interfaces ge-1/2/7 unit 0 family inet address 10.10.4.2/24
set interfaces ge-1/2/8 unit 0 family inet address 10.10.5.2/24
set interfaces xe-2/0/0 unit 0 family inet address 10.10.10.2/24
set interfaces lo0 unit 0 family inet address 10.255.44.44/32
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/7.0
set protocols ospf area 0.0.0.0 interface ge-1/2/8.0
set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim traceoptions file df

```

```

set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface xe-2/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-1/2/7.0 mode bidirectional-sparse-dense
set protocols pim interface ge-1/2/8.0 mode bidirectional-sparse-dense

```

Router R5

```

set interfaces ge-0/0/3 unit 0 family inet address 10.10.12.3/24
set interfaces ge-0/0/4 unit 0 family inet address 10.10.4.3/24
set interfaces ge-0/0/7 unit 0 family inet address 10.10.5.3/24
set interfaces so-1/0/0 unit 0 family inet address 10.10.7.1/30
set interfaces lo0 unit 0 family inet address 10.255.55.55/32
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/0/7.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface so-1/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface ge-0/0/7.0 mode bidirectional-sparse-dense
set protocols pim interface ge-0/0/4.0 mode bidirectional-sparse-dense
set protocols pim interface so-1/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-0/0/3.0 mode bidirectional-sparse-dense

```

Router R6

```

set interfaces xe-0/0/0 unit 0 family inet address 10.10.10.3/24
set interfaces ge-2/0/0 unit 0 family inet address 10.10.13.2/24
set interfaces lo0 unit 0 family inet address 10.255.66.66/32
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24

```

```

set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface fxp0.0 disable
set protocols pim interface xe-0/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-2/0/0.0 mode bidirectional-sparse-dense

```

Router R7

```

set interfaces ge-0/1/5 unit 0 family inet address 10.10.13.3/24
set interfaces ge-0/1/7 unit 0 family inet address 10.10.12.2/24
set interfaces lo0 unit 0 family inet address 10.255.77.77/32
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/1/5.0
set protocols ospf area 0.0.0.0 interface ge-0/1/7.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface ge-0/1/5.0 mode bidirectional-sparse-dense
set protocols pim interface ge-0/1/7.0 mode bidirectional-sparse-dense

```

Router R8

```

set interfaces so-0/0/0 unit 0 family inet address 10.10.7.2/30
set interfaces xe-2/0/0 unit 0 family inet address 10.10.9.2/24
set interfaces lo0 unit 0 family inet address 10.255.88.88/32
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
set protocols ospf area 0.0.0.0 interface so-0/0/0.0
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface xe-2/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface so-0/0/0.0 mode bidirectional-sparse-dense

```

Router R1

Step-by-Step Procedure

To configure Router R1:

1. Configure the router interfaces.

```
[edit interfaces]
user@R1# set ge-0/0/1 unit 0 family inet address 10.10.1.1/24
user@R1# set xe-2/1/0 unit 0 family inet address 10.10.2.1/24
user@R1# set lo0 unit 0 family inet address 10.255.11.11/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface ge-0/0/1.0
user@R1# set interface xe-2/1/0.0
user@R1# set interface lo0.0
user@R1# set interface fxp0.0 disable
```

3. Configure the group-to-RP mappings.

```
[edit protocols pim rp bidirectional]
user@R1# set address 10.10.1.3 group-ranges 224.1.3.0/24
user@R1# set address 10.10.1.3 group-ranges 225.1.3.0/24
user@R1# set address 10.10.13.2 group-ranges 224.1.1.0/24
user@R1# set address 10.10.13.2 group-ranges 225.1.1.0/24
```

The RP represented by IP address 10.10.1.3 is a phantom RP. The 10.10.1.3 address is not assigned to any interface on any of the routers in the topology. It is, however, a reachable address. It is in the subnet between Routers R1 and R2.

The RP represented by address 10.10.13.2 is assigned to the **ge-2/0/0** interface on Router R6.

4. Enable bidirectional PIM on the interfaces.

```
[edit protocols pim]
user@R1# set interface ge-0/0/1.0 mode bidirectional-sparse-dense
user@R1# set interface xe-2/1/0.0 mode bidirectional-sparse-dense
```

5. (Optional) Configure tracing operations for the DF election process.

```
[edit protocols pim]
user@R1# set traceoptions file df
user@R1# set traceoptions flag bidirectional-df-election detail
```

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.

```
user@R1# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.10.1.1/24;
    }
  }
}
xe-2/1/0 {
  unit 0 {
    family inet {
      address 10.10.2.1/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.11.11/32;
    }
  }
}
```

```
user@R1# show protocols
ospf {
  area 0.0.0.0 {
    interface ge-0/0/1.0;
```



```

interface xe-2/1/0.0;
interface lo0.0;
interface fxp0.0 {
    disable;
}
}
}
pim {
    rp {
        bidirectional {
            address 10.10.1.3 { # phantom RP
                group-ranges {
                    224.1.3.0/24;
                    225.1.3.0/24;
                }
            }
            address 10.10.13.2 {
                group-ranges {
                    224.1.1.0/24;
                    225.1.1.0/24;
                }
            }
        }
    }
}
interface ge-0/0/1.0 {
    mode bidirectional-sparse-dense;
}
interface xe-2/1/0.0 {
    mode bidirectional-sparse-dense;
}
traceoptions {
    file df;
    flag bidirectional-df-election detail;
}
}
}

```

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

Repeat the procedure for every Juniper Networks router in the bidirectional PIM network, using the appropriate interface names and addresses for each router.

Verification

IN THIS SECTION

- [Verifying Rendezvous Points | 488](#)
- [Verifying Messages | 489](#)
- [Checking the PIM Join State | 489](#)
- [Displaying the Designated Forwarder | 491](#)
- [Displaying the PIM Interfaces | 491](#)
- [Checking the PIM Neighbors | 492](#)
- [Checking the Route to the Rendezvous Points | 492](#)
- [Verifying Multicast Routes | 493](#)
- [Viewing Multicast Next Hops | 495](#)

Confirm that the configuration is working properly.

Verifying Rendezvous Points

Purpose

Verify the group-to-RP mapping information.

Action

```

user@R1> show pim rps
Instance: PIM.master
Address family INET
RP address      Type      Mode   Holdtime Timeout Groups Group prefixes
10.10.1.3       static   bidir   150     None     2 224.1.3.0/24
                225.1.3.0/24
10.10.13.2      static   bidir   150     None     2 224.1.1.0/24
                225.1.1.0/24

```

Verifying Messages

Purpose

Check the number of DF election messages sent and received, and check bidirectional join and prune error statistics.

Action

```

user@R1> show pim statistics
PIM Message type      Received      Sent  Rx errors
V2 Hello                16           34      0
...
V2 DF Election         18           38      0
...

Global Statistics

...
Rx Bidir Join/Prune on non-Bidir if      0
Rx Bidir Join/Prune on non-DF if        0

```

Checking the PIM Join State

Purpose

Confirm the upstream interface, neighbor, and state information.

Action

```

user@R1> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

```

Upstream neighbor: 10.10.1.2
Upstream state: None
Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)

Group: 224.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-2/1/0.0 (DF Winner)

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)

```

Interface: lo0.0      (DF Winner)
Interface: xe-2/1/0.0 (DF Winner)

```

Meaning

The output shows a (*,G-range) entry for each active bidirectional RP group range. These entries provide a hierarchy from which the individual (*,G) routes inherit RP-derived state (upstream information and accepting interfaces). These entries also provide the control plane basis for the (*, G-range) forwarding routes that implement the sender-only branches of the tree.

Displaying the Designated Forwarder

Purpose

Display RP address information and confirm the DF elected.

Action

```

user@R1> show pim bidirectional df-election
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
  ge-0/0/1.0  (RPL)    DF: none
  lo0.0       (Win)     DF: 10.255.179.246
  xe-2/1/0.0  (Win)     DF: 10.10.2.1

RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
  ge-0/0/1.0  (Lose)    DF: 10.10.1.2
  lo0.0       (Win)     DF: 10.255.179.246
  xe-2/1/0.0  (Lose)    DF: 10.10.2.2

```

Displaying the PIM Interfaces

Purpose

Verify that the PIM interfaces have bidirectional-sparse-dense (**SDB**) mode assigned.

Action

```
user@R1> show pim interfaces
```

```
Instance: PIM.master
```

Stat = Status, V = Version, NbrCnt = Neighbor Count,

S = Sparse, D = Dense, B = Bidirectional,

DR = Designated Router, P2P = Point-to-point link,

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name	Stat	Mode	IP V	State	NbrCnt	JoinCnt(sg/*g)	DR address
ge-0/0/1.0	Up	SDB	4 2	NotDR,Active	1	0/0	10.10.1.2
lo0.0	Up	SDB	4 2	DR,Active	0	9901/100	10.255.179.246
xe-2/1/0.0	Up	SDB	4 2	NotDR,Active	1	0/0	10.10.2.2

Checking the PIM Neighbors

Purpose

Check that the router detects that its neighbors are enabled for bidirectional PIM by verifying that the **B** option is displayed.

Action

```
user@R1> show pim neighbors
```

```
Instance: PIM.master
```

B = Bidirectional Capable, G = Generation Identifier,

H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,

P = Hello Option DR Priority, T = Tracking Bit

Interface	IP V	Mode	Option	Uptime	Neighbor addr
ge-0/0/1.0	4 2		HPLGBT	00:06:46	10.10.1.2
xe-2/1/0.0	4 2		HPLGBT	00:06:46	10.10.2.2

Checking the Route to the Rendezvous Points

Purpose

Check the interface route to the rendezvous points.

Action

```

user@R1> show route 10.10.13.2
inet.0: 56 destinations, 56 routes (55 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.13.0/24      *[OSPF/10] 00:04:35, metric 4
                  > to 10.10.1.2 via ge-0/0/1.0

```

```

user@R1> show route 10.10.1.3
inet.0: 56 destinations, 56 routes (55 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.1.0/24      *[Direct/0] 00:06:25
                  > via ge-0/0/1.0

```

Verifying Multicast Routes

Purpose

Verify the multicast traffic route for each group.

For bidirectional PIM, the **show multicast route extensive** command shows the (*, G/prefix) forwarding routes and the list of interfaces that accept bidirectional PIM traffic.

Action

```

user@R1> show multicast route extensive
Family: INET

Group: 224.0.0.0/4
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
  Downstream interface list:
    ge-0/0/1.0
  Session description: zeroconfaddr
  Statistics: 0 kBps, 0 pps, 0 packets
  Next-hop ID: 2097157

```

Incoming interface list ID: 559
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

Group: 224.1.1.0/24

Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0
Downstream interface list:
 ge-0/0/1.0
Session description: NOB Cross media facilities
Statistics: 0 kBps, 0 pps, 0 packets
Next-hop ID: 2097157
Incoming interface list ID: 579
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

Group: 224.1.3.0/24

Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
 ge-0/0/1.0
Session description: NOB Cross media facilities
Statistics: 0 kBps, 0 pps, 0 packets
Next-hop ID: 2097157
Incoming interface list ID: 556
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

Group: 225.1.1.0/24

Source: *
Incoming interface list:
 lo0.0 ge-0/0/1.0


```

Downstream interface list:
  ge-0/0/1.0
Session description: Unknown
Statistics: 0 kBps, 0 pps, 0 packets
Next-hop ID: 2097157
Incoming interface list ID: 579
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

```

Group: 225.1.3.0/24
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Session description: Unknown
Statistics: 0 kBps, 0 pps, 0 packets
Next-hop ID: 2097157
Incoming interface list ID: 556
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

Meaning

For information about how the incoming and outgoing interface lists are derived, see the forwarding rules in RFC 5015.

Viewing Multicast Next Hops

Purpose

Verify that the correct accepting interfaces are shown in the incoming interface list.

Action

```

user@R1> show multicast next-hops
Family: INET
ID          Refcount KRefCount Downstream interface
2097157     10       5 ge-0/0/1.0

Family: Incoming interface list
ID          Refcount KRefCount Downstream interface
579         5         2 lo0.0
              ge-0/0/1.0
556         5         2 lo0.0
              ge-0/0/1.0
              xe-4/1/0.0
559         3         1 lo0.0
              ge-0/0/1.0
              xe-4/1/0.0

```

Meaning

The nexthop IDs for the outgoing and incoming next hops are referenced directly in the **show multicast route extensive** command.

SEE ALSO

[Understanding Bidirectional PIM | 470](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
13.3	PTX5000 routers do not support nonstop active routing or in-service software upgrade (ISSU) in Junos OS Release 13.3.
12.2	Starting with Release 12.2, Junos OS extends the <i>nonstop active routing</i> PIM support to draft-rosen MVPNs.

Rapidly Detecting Communication Failures with PIM and the BFD Protocol

IN THIS CHAPTER

- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol | 497](#)

Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol

IN THIS SECTION

- [Understanding Bidirectional Forwarding Detection Authentication for PIM | 497](#)
- [Configuring BFD for PIM | 500](#)
- [Configuring BFD Authentication for PIM | 502](#)
- [Example: Configuring BFD Liveness Detection for PIM IPv6 | 506](#)

Understanding Bidirectional Forwarding Detection Authentication for PIM

IN THIS SECTION

- [BFD Authentication Algorithms | 498](#)
- [Security Authentication Keychains | 499](#)
- [Strict Versus Loose Authentication | 499](#)

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over PIM. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

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

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

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

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

secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



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

Security Authentication Keychains

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

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

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



NOTE: Security Authentication Keychain is not supported on SRX Series Firewalls.

Strict Versus Loose Authentication

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

SEE ALSO

[Configuring BFD Authentication for PIM | 291](#)

[Configuring BFD for PIM | 289](#)

[authentication-key-chains](#)

[bfd-liveness-detection](#)

[show bfd session](#)

Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```
[edit protocols pim]
user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the show bfd session command.

SEE ALSO

| *show bfd session*

Configuring BFD Authentication for PIM

IN THIS SECTION

- [Configuring BFD Authentication Parameters | 502](#)
- [Viewing Authentication Information for BFD Sessions | 504](#)

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

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication algorithm
keyed-sha-1
```




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

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

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

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication keychain
bfd-pim
```



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

3. Specify the unique security authentication information for BFD sessions:
 - The matching keychain name as specified in Step 2.
 - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
 - The secret data used to allow access to the session.
 - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret $ABC123$/ start-time
2009-06-14.10:00:00
```



NOTE: Security Authentication Keychain is not supported on SRX Series Firewalls.

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

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication loose-
check
```

5. (Optional) View your configuration by using the `show bfd session detail` or `show bfd session extensive` command.
6. Repeat these steps to configure the other end of the BFD session.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the `show bfd session detail` and `show bfd session extensive` commands.

The following example shows BFD authentication configured for the `ge-0/1/5` interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of `bfd-pim`. The authentication keychain is configured with two keys. Key **1** contains the secret data “`$ABC123/`” and a start time of June 1, 2009, at 9:46:02 AM PST. Key **2** contains the secret data “`$ABC123/`” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface ge-0/1/5 {
  family inet {
    bfd-liveness-detection {
      authentication {
        key-chain bfd-pim;
        algorithm keyed-sha-1;
      }
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-pim {
    key 1 {
      secret "$ABC123/";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$ABC123/";
```

```

        start-time "2009-6-1.15:29:20 -0700";
    }
}
}

```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the `show bfd session detail` command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the `show bfd session extensive` command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.0.2.2	Up	ge-0/1/5.0	0.900	0.300	3

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**
 Session up time 3d 00:34
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated

show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.0.2.2	Up	ge-0/1/5.0	0.900	0.300	3

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**
keychain bfd-pim, algo keyed-sha-1, mode strict
 Session up time 00:04:42
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.300, min slow interval 1.000
 Adaptive async TX interval 0.300, RX interval 0.300
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
 Local discriminator 2, remote discriminator 2

```
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict
```

RELATED DOCUMENTATION

[Understanding Bidirectional Forwarding Detection Authentication for PIM | 497](#)

[Configuring BFD for PIM | 500](#)

[authentication-key-chains](#)

[bfd-liveness-detection](#)

[show bfd session](#)

Example: Configuring BFD Liveness Detection for PIM IPv6

IN THIS SECTION

- [Requirements | 507](#)
- [Overview | 507](#)
- [Configuration | 508](#)
- [Verification | 513](#)

This example shows how to configure Bidirectional Forwarding Detection (BFD) liveness detection for IPv6 interfaces configured for the Protocol Independent Multicast (PIM) topology. BFD is a simple hello mechanism that detects failures in a network.

The following steps are needed to configure BFD liveness detection:

1. Configure the interface.
2. Configure the related security authentication keychain.
3. Specify the BFD authentication algorithm for the PIM protocol.
4. Configure PIM, associating the authentication keychain with the desired protocol.
5. Configure BFD authentication for the routing instance.



NOTE: You must perform these steps on both ends of the BFD session.

Requirements

This example uses the following hardware and software components:

- Two peer routers.
- Junos OS 12.2 or later.

Overview

IN THIS SECTION


- [Topology | 507](#)

In this example, Device R1 and Device R2 are peers. Each router runs PIM, connected over a common medium.

Topology

[Figure 75 on page 507](#) shows the topology used in this example.

Figure 75: BFD Liveness Detection for PIM IPv6 Topology



Assume that the routers initialize. No BFD session is yet established. For each router, PIM informs the BFD process to monitor the IPv6 address of the neighbor that is configured in the routing protocol. Addresses are not learned dynamically and must be configured.

Configure the IPv6 address and BFD liveness detection at the `[edit protocols pim]` hierarchy level for each router.

```
[edit protocols pim]
user@host# set interface interface-name family inet6 bfd-liveness-detection
```

Configure BFD liveness detection for the routing instance at the `[edit routing-instances instance-name protocols pim interface all family inet6]` hierarchy level (here, the *instance-name* is `instance1`):

```
[edit routing-instances instance1 protocols pim]
user@host# set bfd-liveness-detection
```

You will also configure the authentication algorithm and authentication keychain values for BFD.

In a BFD-configured network, when a client launches a BFD session with a peer, BFD begins sending slow, periodic BFD control packets that contain the interval values that you specified when you configured the BFD peers. This is known as the initialization state. BFD does not generate any up or down notifications in this state. When another BFD interface acknowledges the BFD control packets, the session moves into an up state and begins to more rapidly send periodic control packets. If a data path failure occurs and BFD does not receive a control packet within the configured amount of time, the data path is declared down and BFD notifies the BFD client. The BFD client can then perform the necessary actions to reroute traffic. This process can be different for different BFD clients.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 508](#)
- [Procedure | 509](#)
- [Results | 511](#)

CLI Quick Configuration

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

Device R1

```
set interfaces ge-0/1/5 unit 0 description toRouter2
set interfaces ge-0/1/5 unit 0 family inet6
set interfaces ge-0/1/5 unit 0 family inet6 address e80::21b:c0ff:fed5:e4dd
set protocols pim interface ge-0/1/5 family inet6 bfd-liveness-detection authentication
algorithm keyed-sha-1
```

```

set protocols pim interface ge-0/1/5 family inet6 bfd-liveness-detection authentication key-
chain bfd-pim
set routing-instances instance1 protocols pim interface all family inet6 bfd-liveness-detection
authentication algorithm keyed-sha-1
set routing-instances instance1 protocols pim interface all family inet6 bfd-liveness-detection
authentication key-chain bfd-pim
set security authentication key-chain bfd-pim key 1 secret "v"
set security authentication key-chain bfd-pim key 1 start-time "2012-01-01.09:46:02 -0700"
set security authentication key-chain bfd-pim key 2 secret "$ABC123abc123"
set security authentication key-chain bfd-pim key 2 start-time "2012-01-01.15:29:20 -0700"

```

Device R2

```

set interfaces ge-1/1/0 unit 0 description toRouter1
set interfaces ge-1/1/0 unit 0 family inet6 address e80::21b:c0ff:fed5:e5dd
set protocols pim interface ge-1/1/0 family inet6 bfd-liveness-detection authentication
algorithm keyed-sha-1
set protocols pim interface ge-1/1/0 family inet6 bfd-liveness-detection authentication key-
chain bfd-pim
set routing-instances instance1 protocols pim interface all family inet6 bfd-liveness-detection
authentication algorithm keyed-sha-1
set routing-instances instance1 protocols pim interface all family inet6 bfd-liveness-detection
authentication key-chain bfd-pim
set security authentication key-chain bfd-pim key 1 secret "$ABC123abc123"
set security authentication key-chain bfd-pim key 1 start-time "2012-01-01.09:46:02 -0700"
set security authentication key-chain bfd-pim key 2 secret "$ABC123abc123"
set security authentication key-chain bfd-pim key 2 start-time "2012-01-01.15:29:20 -0700"

```

Procedure

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 [Junos OS CLI User Guide](#).

To configure BFD liveness detection for PIM IPv6 interfaces on Device R1:



NOTE: This procedure is for Device R1. Repeat this procedure for Device R2, after modifying the appropriate interface names, addresses, and any other parameters.

1. Configure the interface, using the `inet6` statement to specify that this is an IPv6 address.

```
[edit interfaces]
user@R1# set ge-0/1/5 unit 0 description toRouter2
user@R1# set ge-0/1/5 unit 0 family inet6 address e80::21b:c0ff:fed5:e4dd
```

2. Specify the BFD authentication algorithm and keychain for the PIM protocol.

The keychain is used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes. This keychain name should match the keychain name configured at the `[edit security authentication]` hierarchy level.

```
[edit protocols]
user@R1# set pim interface ge-0/1/5.0 family inet6 bfd-liveness-detection authentication
algorithm keyed-sha-1
user@R1# set pim interface ge-0/1/5 family inet6 bfd-liveness-detection authentication key-
chain bfd-pim
```



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

3. Configure a routing instance (here, **instance1**), specifying BFD authentication and associating the security authentication algorithm and keychain.

```
[edit routing-instances]
user@R1# set instance1 protocols pim interface all family inet6 bfd-liveness-detection
authentication algorithm keyed-sha-1
user@R1# set instance1 protocols pim interface all family inet6 bfd-liveness-detection
authentication key-chain bfd-pim
```

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

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.

- The time at which the authentication key becomes active, in the format **YYYY-MM-DD.hh:mm:ss**.

```
[edit security authentication]
user@R1# set key-chain bfd-pim key 1 secret "$ABC123abc123"
user@R1# set key-chain bfd-pim key 1 start-time "2012-01-01.09:46:02 -0700"
user@R1# set key-chain bfd-pim key 2 secret "$ABC123abc123"
user@R1# set key-chain bfd-pim key 2 start-time "2012-01-01.15:29:20 -0700"
```

Results

Confirm your configuration by issuing the `show interfaces`, `show protocols`, `show routing-instances`, and `show security` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
ge-0/1/5 {
  unit 0 {
    description toRouter2;
    family inet6 {
      address e80::21b:c0ff:fed5:e4dd {
      }
    }
  }
}
```

```
user@R1# show protocols
pim {
  interface ge-0/1/5.0 {
    family inet6;
    bfd-liveness-detection {
      authentication {
        algorithm keyed-sha-1;
        key-chain bfd-pim;
      }
    }
  }
}
```

```
}  
}
```

```
user@R1# show routing-instances  
instance1 {  
  protocols {  
    pim {  
      interface all {  
        family inet6 {  
          bfd-liveness-detection {  
            authentication {  
              algorithm keyed-sha-1;  
              key-chain bfd-pim;  
            }  
          }  
        }  
      }  
    }  
  }  
}
```

```
user@R1# show security  
authentication {  
  key-chain bfd-pim {  
    key 1 {  
      secret "$ABC123abc123";  
      start-time "2012-01-01.09:46:02 -0700";  
    }  
    key 2 {  
      secret "$ABC123abc123";  
      start-time "2012-01-01.15:29:20 -0700";  
    }  
  }  
}
```

Verification

IN THIS SECTION

- [Verifying the BFD Session | 513](#)

Confirm that the configuration is working properly.

Verifying the BFD Session

Purpose

Verify that BFD liveness detection is enabled.

Action

```
user@R1# run show pim neighbors detail
```

```
Instance: PIM.master
```

```
Interface: ge-0/1/5.0
```

```
Address: fe80::21b:c0ff:fed5:e4dd, IPv6, PIM v2, Mode: Sparse, sg Join Count: 0, tsg
Join Count: 0
```

```
Hello Option Holdtime: 65535 seconds
```

```
Hello Option DR Priority: 1
```

```
Hello Option Generation ID: 1417610277
```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Join Suppression supported
```

```
Address: fe80::21b:c0ff:fedc:28dd, IPv6, PIM v2, sg Join Count: 0, tsg Join Count: 0
```

```
Secondary address: beef::2
```

```
BFD: Enabled, Operational state: Up
```

```
Hello Option Holdtime: 105 seconds 80 remaining
```

```
Hello Option DR Priority: 1
```

```
Hello Option Generation ID: 1648636754
```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Join Suppression supported
```

Meaning

The display from the `show pim neighbors detail` command shows **BFD: Enabled, Operational state: Up**, indicating that BFD is operating between the two PIM neighbors. For additional information about the BFD session (including the session ID number), use the `show bfd session extensive` command.

SEE ALSO

[authentication-key-chains](#)

bfd-liveness-detection (Protocols PIM)

show bfd session

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
9.6	Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over PIM. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.
9.6	Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported.

RELATED DOCUMENTATION

[Configuring Basic PIM Settings](#)

[Example: Configuring BFD for BGP](#)

[Example: Configuring BFD Authentication for BGP](#)

Configuring PIM Options

IN THIS CHAPTER

- [Example: Configuring Nonstop Active Routing for PIM | 515](#)
- [Configuring PIM-to-IGMP and PIM-to-MLD Message Translation | 534](#)

Example: Configuring Nonstop Active Routing for PIM

IN THIS SECTION

- [Understanding Nonstop Active Routing for PIM | 515](#)
- [Example: Configuring Nonstop Active Routing with PIM | 516](#)
- [Configuring PIM Sparse Mode Graceful Restart | 532](#)

Understanding Nonstop Active Routing for PIM

Nonstop active routing configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. When nonstop active routing 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

The PIM control state is maintained on the backup Routing Engine by the replication of state information from the primary to the backup Routing Engine and having the backup Routing Engine react to route installation and modification in the **[instance].inet.1** routing table on the primary Routing Engine. The backup Routing Engine does not send or receive PIM protocol packets directly. In addition, the backup Routing Engine uses the dynamic interfaces created by the primary Routing Engine. These dynamic interfaces include PIM encapsulation, de-encapsulation, and multicast tunnel interfaces.



NOTE: The **clear pim join**, **clear pim register**, and **clear pim statistics** operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

To enable nonstop active routing for PIM (in addition to the PIM configuration on the primary Routing Engine), you must include the following statements at the **[edit]** hierarchy level:

- **chassis redundancy graceful-switchover**
- **routing-options nonstop-routing**
- **system commit synchronize**

Example: Configuring Nonstop Active Routing with PIM

IN THIS SECTION

- [Requirements | 516](#)
- [Overview | 517](#)
- [Configuration | 518](#)
- [Verification | 531](#)

This example shows how to configure nonstop active routing for PIM-based multicast IPv4 and IPv6 traffic.

Requirements

For nonstop active routing for PIM-based multicast traffic to work with IPv6, the routing device must be running Junos OS Release 10.4 or above.

Before you begin:

- Configure the router interfaces. See the *Network Interfaces Configuration Guide*.

- Configure an interior gateway protocol or static routing. See the *Routing Protocols Configuration Guide*.
- Configure a multicast group membership protocol (IGMP or MLD). See "[Understanding IGMP](#)" on [page 27](#) and "[Understanding MLD](#)" on [page 59](#).

Overview

IN THIS SECTION

- [Topology](#) | **518**

Junos OS supports nonstop active routing 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 (use the `advertise-from-main-vpn-tables` option at the `[edit protocols bgp]` hierarchy level, to synchronize MVPN routes, `cmcast`, `provider-tunnel` and forwarding information between the primary and the backup Routing Engines).
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies

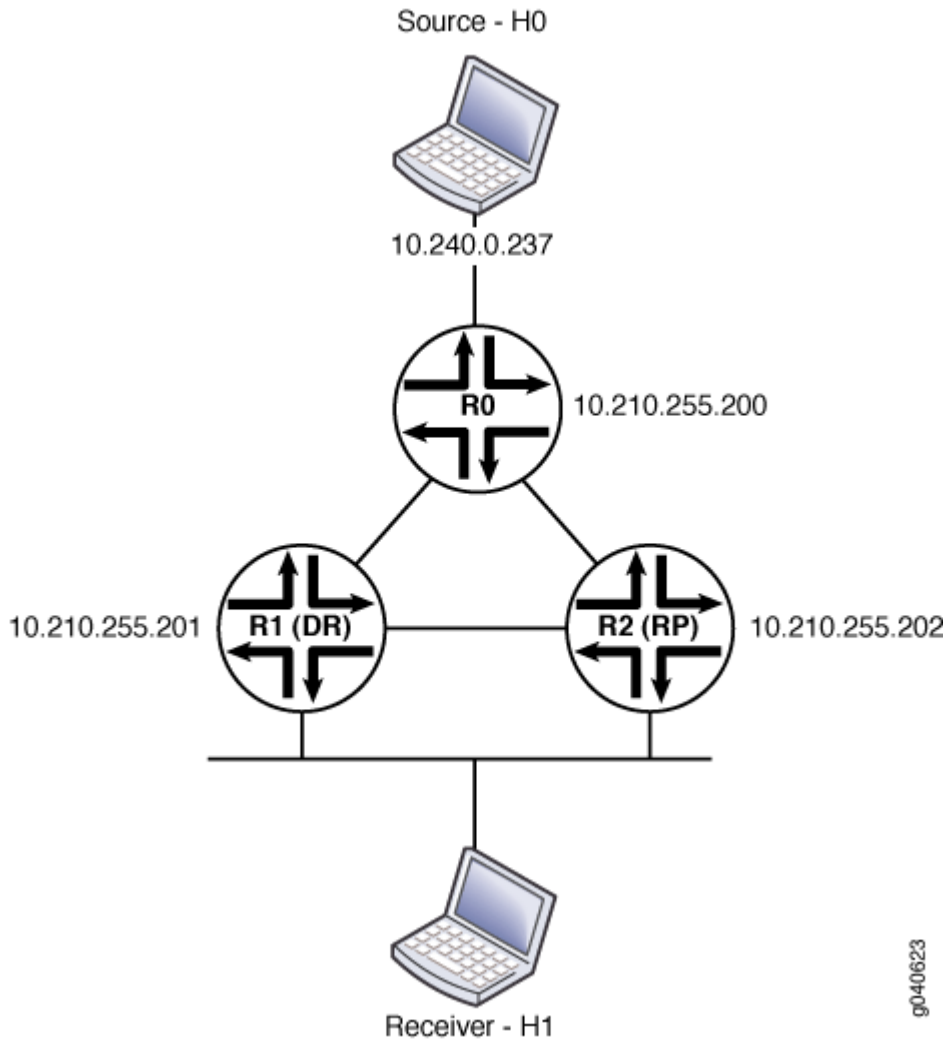
In Junos OS release 13.3, multicast VPNs are not supported with nonstop active routing. Policy-based features (such as neighbor policy, join policy, BSR policy, scope policy, flow maps, and RPF check policy) are not supported with nonstop active routing.

This example uses static RP. The interfaces are configured to receive both IPv4 and IPv6 traffic. R2 provides RP services as the local RP. Note that nonstop active routing is not supported on the RP router. The configuration shown in this example is on R1.

Topology

Figure 76 on page 518 shows the topology used in this example.

Figure 76: Nonstop Active Routing in PIM Domain



Configuration

IN THIS SECTION

- CLI Quick Configuration | 519
- Procedure | 521

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.

R1

```
set system syslog archive size 10m
set system syslog file messages any info
set system commit synchronize
set chassis redundancy graceful-switchover
set interfaces traceoptions file dcd-trace
set interfaces traceoptions file size 10m
set interfaces traceoptions file files 10
set interfaces traceoptions flag all
set interfaces so-0/0/1 unit 0 description "to R0 so-0/0/1.0"
set interfaces so-0/0/1 unit 0 family inet address 10.210.1.2/30
set interfaces so-0/0/1 unit 0 family inet6 address FDCA:9E34:50CE:0001::2/126
set interfaces fe-0/1/3 unit 0 description "to R2 fe-0/1/3.0"
set interfaces fe-0/1/3 unit 0 family inet address 10.210.12.1/30
set interfaces fe-0/1/3 unit 0 family inet6 address FDCA:9E34:50CE:0012::1/126
set interfaces fe-1/1/0 unit 0 description "to H1"
set interfaces fe-1/1/0 unit 0 family inet address 10.240.0.250/30
set interfaces fe-1/1/0 unit 0 family inet6 address ::10.240.0.250/126
set interfaces lo0 unit 0 description "R1 Loopback"
set interfaces lo0 unit 0 family inet address 10.210.255.201/32 primary
set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.1025.5201.00
set interfaces lo0 unit 0 family inet6 address abcd::10:210:255:201/128
set protocols ospf traceoptions file r1-nsr-ospf2
set protocols ospf traceoptions file size 10m
set protocols ospf traceoptions file files 10
set protocols ospf traceoptions file world-readable
set protocols ospf traceoptions flag error
set protocols ospf traceoptions flag lsa-update detail
set protocols ospf traceoptions flag flooding detail
set protocols ospf traceoptions flag lsa-request detail
```

```
set protocols ospf traceoptions flag state detail
set protocols ospf traceoptions flag event detail
set protocols ospf traceoptions flag hello detail
set protocols ospf traceoptions flag nsr-synchronization detail
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface so-0/0/1.0 metric 100
set protocols ospf area 0.0.0.0 interface fe-0/1/3.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface fe-1/1/0.0 passive
set protocols ospf3 traceoptions file r1-nsr-ospf3
set protocols ospf3 traceoptions file size 10m
set protocols ospf3 traceoptions file world-readable
set protocols ospf3 traceoptions flag lsa-update detail
set protocols ospf3 traceoptions flag flooding detail
set protocols ospf3 traceoptions flag lsa-request detail
set protocols ospf3 traceoptions flag state detail
set protocols ospf3 traceoptions flag event detail
set protocols ospf3 traceoptions flag hello detail
set protocols ospf3 traceoptions flag nsr-synchronization detail
set protocols ospf3 area 0.0.0.0 interface fe-1/1/0.0 passive
set protocols ospf3 area 0.0.0.0 interface fe-1/1/0.0 metric 1
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface so-0/0/1.0 metric 1
set protocols ospf3 area 0.0.0.0 interface fe-0/1/3.0 metric 1
set protocols pim traceoptions file r1-nsr-pim
set protocols pim traceoptions file size 10m
set protocols pim traceoptions file files 10
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag mdt detail
set protocols pim traceoptions flag rp detail
set protocols pim traceoptions flag register detail
set protocols pim traceoptions flag packets detail
set protocols pim traceoptions flag autorp detail
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag hello detail
set protocols pim traceoptions flag assert detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag state detail
set protocols pim traceoptions flag nsr-synchronization
set protocols pim rp static address 10.210.255.202
set protocols pim rp static address abcd::10:210:255:202
set protocols pim interface lo0.0
```

```
set protocols pim interface fe-0/1/3.0 mode sparse
set protocols pim interface fe-0/1/3.0 version 2
set protocols pim interface so-0/0/1.0 mode sparse
set protocols pim interface so-0/0/1.0 version 2
set protocols pim interface fe-1/1/0.0 mode sparse
set protocols pim interface fe-1/1/0.0 version 2
set policy-options policy-statement load-balance then load-balance per-packet
set routing-options nonstop-routing
set routing-options router-id 10.210.255.201
set routing-options forwarding-table export load-balance
set routing-options forwarding-table traceoptions file r1-nsr-krt
set routing-options forwarding-table traceoptions file size 10m
set routing-options forwarding-table traceoptions file world-readable
set routing-options forwarding-table traceoptions flag queue
set routing-options forwarding-table traceoptions flag route
set routing-options forwarding-table traceoptions flag routes
set routing-options forwarding-table traceoptions flag synchronous
set routing-options forwarding-table traceoptions flag state
set routing-options forwarding-table traceoptions flag asynchronous
set routing-options forwarding-table traceoptions flag consistency-checking
set routing-options traceoptions file r1-nsr-sync
set routing-options traceoptions file size 10m
set routing-options traceoptions flag nsr-synchronization
set routing-options traceoptions flag commit-synchronize
```

Procedure

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 [Junos OS CLI User Guide](#).

To configure nonstop active routing on R1:

1. Synchronize the Routing Engines.

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

```

user@host# set commit synchronize
user@host# exit

```

2. Enable graceful Routing Engine switchover.

```

[edit]
user@host# set chassis redundancy graceful-switchover

```

3. Configure R1's interfaces.

```

[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set so-0/0/1 unit 0 description "to R0 so-0/0/1.0"
user@host# set so-0/0/1 unit 0 family inet address 10.210.1.2/30
user@host# set so-0/0/1 unit 0 family inet6 address FDCA:9E34:50CE:0001::2/126
user@host# set fe-0/1/3 unit 0 description "to R2 fe-0/1/3.0"
user@host# set fe-0/1/3 unit 0 family inet address 10.210.12.1/30
user@host# set fe-0/1/3 unit 0 family inet6 address FDCA:9E34:50CE:0012::1/126
user@host# set fe-1/1/0 unit 0 description "to H1"
user@host# set fe-1/1/0 unit 0 family inet address 10.240.0.250/30
user@host# set fe-1/1/0 unit 0 family inet6 address ::10.240.0.250/126
user@host# set lo0 unit 0 description "R1 Loopback"
user@host# set lo0 unit 0 family inet address 10.210.255.201/32 primary
user@host# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.1025.5201.00
user@host# set lo0 unit 0 family inet6 address abcd::10:210:255:201/128
user@host# exit

```

4. Configure OSPF for IPv4 on R1.

```

[edit]
user@host# edit protocols ospf
[edit protocols ospf]
user@host# set traffic-engineering
user@host# set area 0.0.0.0 interface so-0/0/1.0 metric 100
user@host# set area 0.0.0.0 interface fe-0/1/3.0 metric 100
user@host# set area 0.0.0.0 interface lo0.0 passive

```

```

user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface fe-1/1/0.0 passive

```

5. Configure OSPF for IPv6 on R1.

```

[edit]
user@host# edit protocols ospf3
[edit protocols ospf3]
user@host# set area 0.0.0.0 interface fe-1/1/0.0 passive
user@host# set area 0.0.0.0 interface fe-1/1/0.0 metric 1
user@host# set area 0.0.0.0 interface lo0.0 passive
user@host# set area 0.0.0.0 interface so-0/0/1.0 metric 1
user@host# set area 0.0.0.0 interface fe-0/1/3.0 metric 1

```

6. Configure PIM on R1. The PIM static address points to the RP router (R2).

```

[edit]
user@host# edit
[edit protocols pim]
user@host# set protocols pim rpstatic address 10.210.255.202
user@host# set protocols pim rp static address abcd::10:210:255:202
user@host# set protocols pim interface (Protocols PIM) lo0.0
user@host# set protocols pim interface fe-0/1/3.0 mode sparse
user@host# set protocols pim interface fe-0/1/3.0 version 2
user@host# set protocols pim interface so-0/0/1.0 mode sparse
user@host# set protocols pim interface so-0/0/1.0 version 2
user@host# set protocols pim interface fe-1/1/0.0 mode sparse
user@host# set protocols pim interface fe-1/1/0.0 version 2

```

7. Configure per-packet load balancing on R1.

```

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

```

8. Apply the load-balance policy on R1.

```
[edit]
user@host# set routing-options forwarding-table export load-balance
```

9. Configure nonstop routing on R1.

```
[edit]
user@host# set routing-options nonstop-routing
user@host# set routing-options router-id 10.210.255.201
```

Step-by-Step Procedure

For troubleshooting, configure system log and tracing operations.

1. Enable system log messages.

```
[edit]
user@host# set system syslog archive size 10m
user@host# set system syslog file messages any info
```

2. Trace interface operations.

```
[edit]
user@host# set interfaces traceoptions file dcd-trace
user@host# set interfaces traceoptions file size 10m
user@host# set interfaces traceoptions file files 10
user@host# set interfaces traceoptions flag all
```

3. Trace IGP operations for IPv4.

```
[edit]
user@host# set protocols ospf traceoptions file r1-nsr-ospf2
user@host# set protocols ospf traceoptions file size 10m
user@host# set protocols ospf traceoptions file files 10
user@host# set protocols ospf traceoptions file world-readable
user@host# set protocols ospf traceoptions flag error
user@host# set protocols ospf traceoptions flag lsa-update detail
```

```
user@host# set protocols ospf traceoptions flag flooding detail
user@host# set protocols ospf traceoptions flag lsa-request detail
user@host# set protocols ospf traceoptions flag state detail
user@host# set protocols ospf traceoptions flag event detail
user@host# set protocols ospf traceoptions flag hello detail
user@host# set protocols ospf traceoptions flag nsr-synchronization detail
```

4. Trace IGP operations for IPv6.

```
[edit]
user@host# set protocols ospf3 traceoptions file r1-nsr-ospf3
user@host# set protocols ospf3 traceoptions file size 10m
user@host# set protocols ospf3 traceoptions file world-readable
user@host# set protocols ospf3 traceoptions flag lsa-update detail
user@host# set protocols ospf3 traceoptions flag flooding detail
user@host# set protocols ospf3 traceoptions flag lsa-request detail
user@host# set protocols ospf3 traceoptions flag state detail
user@host# set protocols ospf3 traceoptions flag event detail
user@host# set protocols ospf3 traceoptions flag hello detail
user@host# set protocols ospf3 traceoptions flag nsr-synchronization detail
```

5. Trace PIM operations.

```
[edit]
user@host# set protocols pim traceoptions file r1-nsr-pim
user@host# set protocols pim traceoptions file size 10m
user@host# set protocols pim traceoptions file files 10
user@host# set protocols pim traceoptions file world-readable
user@host# set protocols pim traceoptions flag mdt detail
user@host# set protocols pim traceoptions flag rp detail
user@host# set protocols pim traceoptions flag register detail
user@host# set protocols pim traceoptions flag packets detail
user@host# set protocols pim traceoptions flag autorp detail
user@host# set protocols pim traceoptions flag join detail
user@host# set protocols pim traceoptions flag hello detail
user@host# set protocols pim traceoptions flag assert detail
user@host# set protocols pim traceoptions flag normal detail
user@host# set protocols pim traceoptions flag state detail
user@host# set protocols pim traceoptions flag nsr-synchronization
```

6. Trace all routing protocol functionality.

```
[edit]
user@host# set routing-options traceoptions file r1-nsr-sync
user@host# set routing-options traceoptions file size 10m
user@host# set routing-options traceoptions flag nsr-synchronization
user@host# set routing-options traceoptions flag commit-synchronize
```

7. Trace forwarding table operations.

```
[edit]
user@host# set routing-options forwarding-table traceoptions file r1-nsr-krt
user@host# set routing-options forwarding-table traceoptions file size 10m
user@host# set routing-options forwarding-table traceoptions file world-readable
user@host# set routing-options forwarding-table traceoptions flag queue
user@host# set routing-options forwarding-table traceoptions flag route
user@host# set routing-options forwarding-table traceoptions flag routes
user@host# set routing-options forwarding-table traceoptions flag synchronous
user@host# set routing-options forwarding-table traceoptions flag state
user@host# set routing-options forwarding-table traceoptions flag asynchronous
user@host# set routing-options forwarding-table traceoptions flag consistency-checking
```

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

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

Results

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

```
user@host# show chassis
redundancy {
```



```
    graceful-switchover;  
}
```

```
user@host# show interfaces  
traceoptions {  
  file dcd-trace size 10m files 10;  
  flag all;  
}  
so-0/0/1 {  
  unit 0 {  
    description "to R0 so-0/0/1.0";  
    family inet {  
      address 10.210.1.2/30;  
    }  
    family inet6 {  
      address FDCA:9E34:50CE:0001::2/126;  
    }  
  }  
}  
fe-0/1/3 {  
  unit 0 {  
    description "to R2 fe-0/1/3.0";  
    family inet {  
      address 10.210.12.1/30;  
    }  
    family inet6 {  
      address FDCA:9E34:50CE:0012::1/126;  
    }  
  }  
}  
fe-1/1/0 {  
  unit 0 {  
    description "to H1";  
    family inet {  
      address 10.240.0.250/30;  
    }  
    family inet6 {  
      address ::10.240.0.250/126;  
    }  
  }  
}
```

```

lo0 {
  unit 0 {
    description "R1 Loopback";
    family inet {
      address 10.210.255.201/32 {
        primary;
      }
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.1025.5201.00;
    }
    family inet6 {
      address abcd::10:210:255:201/128;
    }
  }
}

```

```

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

```

```

user@host# show protocols
ospf {
  traceoptions {
    file r1-nsr-ospf2 size 10m files 10 world-readable;
    flag error;
    flag lsa-update detail;
    flag flooding detail;
    flag lsa-request detail;
    flag state detail;
    flag event detail;
    flag hello detail;
    flag nsr-synchronization detail;
  }
  traffic-engineering;
  area 0.0.0.0 {
    interface so-0/0/1.0 {

```

```
        metric 100;
    }
    interface fe-0/1/3.0 {
        metric 100;
    }
    interface lo0.0 {
        passive;
    }
    interface fxp0.0 {
        disable;
    }
    interface fe-1/1/0.0 {
        passive;
    }
}
}
ospf3 {
    traceoptions {
        file r1-nsr-ospf3 size 10m world-readable;
        flag lsa-update detail;
        flag flooding detail;
        flag lsa-request detail;
        flag state detail;
        flag event detail;
        flag hello detail;
        flag nsr-synchronization detail;
    }
    area 0.0.0.0 {
        interface fe-1/1/0.0 {
            passive;
            metric 1;
        }
        interface lo0.0 {
            passive;
        }
        interface so-0/0/1.0 {
            metric 1;
        }
        interface fe-0/1/3.0 {
            metric 1;
        }
    }
}
```

```
pim {
  traceoptions {
    file r1-nsr-pim size 10m files 10 world-readable;
    flag mdt detail;
    flag rp detail;
    flag register detail;
    flag packets detail;
    flag autorp detail;
    flag join detail;
    flag hello detail;
    flag assert detail;
    flag normal detail;
    flag state detail;
    flag nsr-synchronization;
  }
  rp {
    static {
      address 10.210.255.202;
      address abcd::10:210:255:202;
    }
  }
  interface lo0.0;
  interface fe-0/1/3.0 {
    mode sparse;
    version 2;
  }
  interface so-0/0/1.0 {
    mode sparse;
    version 2;
  }
  interface fe-1/1/0.0 {
    mode sparse;
    version 2;
  }
}
```

```
user@host# show routing-options
traceoptions {
  file r1-nsr-sync size 10m;
  flag nsr-synchronization;
  flag commit-synchronize;
```

```
}
nonstop-routing;
router-id 10.210.255.201;
forwarding-table {
    traceoptions {
        file r1-nsr-krt size 10m world-readable;
        flag queue;
        flag route;
        flag routes;
        flag synchronous;
        flag state;
        flag asynchronous;
        flag consistency-checking;
    }
    export load-balance;
}
```

```
user@host# show system
syslog {
    archive size 10m;
    file messages {
        any info;
    }
}
commit synchronize;
```

Verification

To verify the configuration, run the following commands:

- **show pim join extensive**
- **show pim neighbors inet detail**
- **show pim neighbors inet6 detail**
- **show pim rps inet detail**
- **show pim rps inet6 detail**
- **show multicast route inet extensive**
- **show multicast route inet6 extensive**

- `show route table inet.1 detail`
- `show route table inet6.1 detail`

Configuring PIM Sparse Mode Graceful Restart

You can configure PIM sparse mode to continue to forward existing multicast packet streams during a routing process failure and restart. Only PIM sparse mode can be configured this way. The routing platform does not forward multicast packets for protocols other than PIM during graceful restart, because all other multicast protocols must restart after a routing process failure. If you configure PIM sparse-dense mode, only sparse multicast groups benefit from a graceful restart.

The routing platform does not forward new streams until after the restart is complete. After restart, the routing platform refreshes the forwarding state with any updates that were received from neighbors during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but it does not apply the changes to the forwarding table until after the restart.

When PIM sparse mode is enabled, the routing platform generates a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the Internet draft [draft-ietf-pim-sm-v2-new-10.txt](#). When a routing platform receives PIM hello messages containing generation identifiers on a point-to-point interface, the Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a routing platform with PIM sparse mode restarts, it creates a new generation identifier and sends it to neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase is complete when the restart duration timer expires.

Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast RPF checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

You can configure graceful restart globally or for a routing instance. This example shows how to configure graceful restart globally.

To configure graceful restart for PIM sparse mode:

1. Enable graceful restart.

```
[edit protocols pim]
user@host# set graceful-restart
```

2. (Optional) Configure the amount of time the routing device waits (in seconds) to complete PIM sparse mode graceful restart. By default, the router allows 60 seconds. The range is from 30 through 300 seconds. After this restart time, the Routing Engine resumes normal multicast operation.

```
[edit protocols pim graceful-restart]
user@host# set restart-duration 120
```

3. Monitor the operation of PIM graceful restart by running the `show pim neighbors` command. In the command output, look for the **G** flag in the **Option** field. The **G** flag stands for generation identifier. Also run the `show task replication` command to verify the status of GRES and NSR.

SEE ALSO

| [Junos OS High Availability User Guide](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
13.3	In Junos OS release 13.3, multicast VPNs are not supported with nonstop active routing. Policy-based features (such as neighbor policy, join policy, BSR policy, scope policy, flow maps, and RPF check policy) are not supported with nonstop active routing.
10.4	For nonstop active routing for PIM-based multicast traffic to work with IPv6, the routing device must be running Junos OS Release 10.4 or above.

RELATED DOCUMENTATION

| [Configuring Basic PIM Settings](#)

Configuring PIM-to-IGMP and PIM-to-MLD Message Translation

IN THIS SECTION

- [Understanding PIM-to-IGMP and PIM-to-MLD Message Translation | 534](#)
- [Configuring PIM-to-IGMP Message Translation | 535](#)
- [Configuring PIM-to-MLD Message Translation | 537](#)

Understanding PIM-to-IGMP and PIM-to-MLD Message Translation

Routing devices can translate Protocol Independent Multicast (PIM) join and prune messages into corresponding Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) report or leave messages. You can use this feature to forward multicast traffic across PIM domains in certain network topologies.

In some network configurations, customers are unable to run PIM between the customer edge-facing PIM domain and the core-facing PIM domain, even though PIM is running in sparse mode within each of these domains. Because PIM is not running between the domains, customers with this configuration cannot use PIM to forward multicast traffic across the domains. Instead, they might want to use IGMP to forward IPv4 multicast traffic, or MLD to forward IPv6 multicast traffic across the domains.

To enable the use of IGMP or MLD to forward multicast traffic across the PIM domains in such topologies, you can configure the rendezvous point (RP) router that resides between the edge domain and core domain to translate PIM join or prune messages received from PIM neighbors on downstream interfaces into corresponding IGMP or MLD report or leave messages. The router then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP router. As a result, this feature is sometimes referred to as *PIM-to-IGMP proxy* or *PIM-to-MLD proxy*.

To configure the RP router to translate PIM join or prune messages into IGMP report or leave messages, include the `pim-to-igmp-proxy` statement at the `[edit routing-options multicast]` hierarchy level. Similarly, to configure the RP router to translate PIM join or prune messages into MLD report or leave messages, include the `pim-to-ml-d-proxy` statement at the `[edit routing-options multicast]` hierarchy level. As part of the configuration, you must specify the full name of at least one, but not more than two, upstream interfaces on which to enable the PIM-to-IGMP proxy or PIM-to-MLD proxy feature.

The following guidelines apply when you configure PIM-to-IGMP or PIM-to-MLD message translation:

- Make sure that the router connecting the PIM edge domain and the PIM core domain is the static or elected RP router.

- Make sure that the RP router is using the PIM sparse mode (PIM-SM) multicast routing protocol.
- When you configure an upstream interface, use the full *logical interface* specification (for example, **ge-0/0/1.0**) and not just the physical interface specification (**ge-0/0/1**).
- When you configure two upstream interfaces, the RP router transmits the same IGMP or MLD report messages and multicast traffic on both upstream interfaces. As a result, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.
- The router transmits IGMP or MLD report messages on one or both upstream interfaces only for the first PIM join message that it receives among all of the downstream interfaces. Similarly, the router transmits IGMP or MLD leave messages on one or both upstream interfaces only if it receives a PIM prune message for the last downstream interface.
- Upstream interfaces support both local sources and remote sources.
- Multicast traffic received from an upstream interface is accepted as if it came from a host.

SEE ALSO

[Understanding PIM Sparse Mode | 307](#)

[Enabling PIM Sparse Mode | 316](#)

Configuring PIM-to-IGMP Message Translation

You can configure the rendezvous point (RP) routing device to translate PIM join or prune messages into corresponding IGMP report or leave messages. To do so, include the `pim-to-igmp-proxy` statement at the `[edit routing-options multicast]` hierarchy level:

```
[edit routing-options multicast]
pim-to-igmp-proxy {
  upstream-interface [ interface-names ];
}
```

Enabling the routing device to perform PIM-to-IGMP message translation, also referred to as *PIM-to-IGMP proxy*, is useful when you want to use IGMP to forward IPv4 multicast traffic between a PIM sparse mode edge domain and a PIM sparse mode core domain in certain network topologies.

Before you begin configuring PIM-to-IGMP message translation:

- Make sure that the routing device connecting the PIM edge domain and that the PIM core domain is the static or elected RP routing device.

- Make sure that the PIM sparse mode (PIM-SM) routing protocol is running on the RP routing device.
- If you plan to configure two upstream interfaces, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain. Because the RP router transmits the same IGMP messages and multicast traffic on both upstream interfaces, you need to run RPF to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

To configure the RP routing device to translate PIM join or prune messages into corresponding IGMP report or leave messages:

1. Include the `pim-to-igmp-proxy` statement, specifying the names of one or two logical interfaces to function as the upstream interfaces on which the routing device transmits IGMP report or leave messages.

The following example configures PIM-to-IGMP message translation on a single upstream interface, **ge-0/1/0.1**.

```
[edit routing-options multicast]
user@host# set pim-to-igmp-proxy upstream-interface ge-0/1/0.1
```

The following example configures PIM-to-IGMP message translation on two upstream interfaces, **ge-0/1/0.1** and **ge-0/1/0.2**. You must include the logical interface names within square brackets (`[]`) when you configure a set of two upstream interfaces.

```
[edit routing-options multicast]
user@host# set pim-to-igmp-proxy upstream-interface [ge-0/1/0.1 ge-0/1/0.2]
```

2. Use the `show multicast pim-to-igmp-proxy` command to display the PIM-to-IGMP proxy state (enabled or disabled) and the name or names of the configured upstream interfaces.

```
user@host# run show multicast pim-to-igmp-proxy
Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

SEE ALSO

pim-to-igmp-proxy
upstream-interface

Configuring PIM-to-MLD Message Translation

You can configure the rendezvous point (RP) routing device to translate PIM join or prune messages into corresponding MLD report or leave messages. To do so, include the `pim-to-ml-d-proxy` statement at the `[edit routing-options multicast]` hierarchy level:

```
[edit routing-options multicast]
pim-to-ml-d-proxy {
  upstream-interface [ interface-names ];
}
```

Enabling the routing device to perform PIM-to-MLD message translation, also referred to as *PIM-to-MLD proxy*, is useful when you want to use MLD to forward IPv6 multicast traffic between a PIM sparse mode edge domain and a PIM sparse mode core domain in certain network topologies.

Before you begin configuring PIM-to-MLD message translation:

- Make sure that the routing device connecting the PIM edge domain and that the PIM core domain is the static or elected RP routing device.
- Make sure that the PIM sparse mode (PIM-SM) routing protocol is running on the RP routing device.
- If you plan to configure two upstream interfaces, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain. Because the RP routing device transmits the same MLD messages and multicast traffic on both upstream interfaces, you need to run RPF to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

To configure the RP routing device to translate PIM join or prune messages into corresponding MLD report or leave messages:

1. Include the `pim-to-ml-d-proxy` statement, specifying the names of one or two logical interfaces to function as the upstream interfaces on which the router transmits MLD report or leave messages. The following example configures PIM-to-MLD message translation on a single upstream interface, **ge-0/5/0.1**.

```
[edit routing-options multicast]
user@host# set pim-to-ml-d-proxy upstream-interface ge-0/5/0.1
```

The following example configures PIM-to-MLD message translation on two upstream interfaces, **ge-0/5/0.1** and **ge-0/5/0.2**. You must include the logical interface names within square brackets ([]) when you configure a set of two upstream interfaces.

```
[edit routing-options multicast]
user@host# set pim-to-mld-proxy upstream-interface [ge-0/5/0.1 ge-0/5/0.2]
```

2. Use the `show multicast pim-to-mld-proxy` command to display the PIM-to-MLD proxy state (enabled or disabled) and the name or names of the configured upstream interfaces.

```
user@host# run show multicast pim-to-mld-proxy
Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

SEE ALSO

[*pim-to-mld-proxy*](#)

[*upstream-interface*](#)

RELATED DOCUMENTATION

[Configuring IGMP | 25](#)

[Configuring MLD | 59](#)

Verifying PIM Configurations

IN THIS CHAPTER

- Verifying the PIM Mode and Interface Configuration | 539
- Verifying the PIM RP Configuration | 540
- Verifying the RPF Routing Table Configuration | 541

Verifying the PIM Mode and Interface Configuration

IN THIS SECTION

- Purpose | 539
- Action | 539
- Meaning | 540

Purpose

Verify that PIM sparse mode is configured on all applicable interfaces.

Action

From the CLI, enter the `show pim interfaces` command.

Sample Output

command-name

```

user@host> show pim interfaces
Instance: PIM.master
Name           Stat Mode      IP V State Count DR address
lo0.0          Up   Sparse      4 2 DR      0 127.0.0.1
pime.32769     Up   Sparse      4 2 P2P     0

```

Meaning

The output shows a list of the interfaces that are configured for PIM. Verify the following information:

- Each interface on which PIM is enabled is listed.
- The network management interface, either **ge-0/0/0** or **fe-0/0/0**, is *not* listed.
- Under **Mode**, the word **Sparse** appears.

Verifying the PIM RP Configuration

IN THIS SECTION

- [Purpose | 540](#)
- [Action | 540](#)
- [Meaning | 541](#)

Purpose

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

Action

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

Sample Output

command-name

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type      Holdtime Timeout Active groups Group prefixes
192.168.14.27  static    0        None      2 224.0.0.0/4

```

Meaning

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

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

Verifying the RPF Routing Table Configuration

IN THIS SECTION

- [Purpose | 541](#)
- [Action | 541](#)
- [Meaning | 542](#)

Purpose

Verify that the PIM RPF routing table is configured correctly.

Action

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

Sample Output

command-name

```
user@host> show multicast rpf
Multicast RPF table: inet.0 , 2 entries...
```

Meaning

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

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

4

PART

Configuring Multicast Routing Protocols

Connecting Routing Domains Using MSDP | 544

Handling Session Announcements with SAP and SDP | 573

Facilitating Multicast Delivery Across Unicast-Only Networks with AMT | 577

Routing Content to Densely Clustered Receivers with DVMRP | 594

Connecting Routing Domains Using MSDP

IN THIS CHAPTER

- [Examples: Configuring MSDP | 544](#)
- [Configuring Multiple Instances of MSDP | 571](#)

Examples: Configuring MSDP

IN THIS SECTION

- [Understanding MSDP | 544](#)
- [Configuring MSDP | 546](#)
- [Example: Configuring MSDP in a Routing Instance | 548](#)
- [Configuring the Interface to Accept Traffic from a Remote Source | 557](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups | 559](#)
- [Tracing MSDP Protocol Traffic | 566](#)
- [Disabling MSDP | 569](#)
- [Example: Configuring MSDP | 570](#)

Understanding MSDP

The Multicast Source Discovery Protocol (MSDP) is used to connect multicast routing domains. It typically runs on the same router as the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). Each MSDP router establishes adjacencies with internal and external MSDP peers similar to the way BGP establishes peers. These peer routers inform each other about active sources within the domain. When they detect active sources, the routers can send PIM sparse-mode explicit join messages to the active source.

The peer with the higher IP address passively listens to a well-known port number and waits for the side with the lower IP address to establish a Transmission Control Protocol (TCP) connection. When a PIM sparse-mode RP that is running MSDP becomes aware of a new local source, it sends source-active type, length, and values (TLVs) to its MSDP peers. When a source-active TLV is received, a peer-reverse-path-forwarding (peer-RPF) check (not the same as a multicast RPF check) is done to make sure that this peer is in the path that leads back to the originating RP. If not, the source-active TLV is dropped. This TLV is counted as a “rejected” source-active message.

The MSDP peer-RPF check is different from the normal RPF checks done by non-MSDP multicast routers. The goal of the peer-RPF check is to stop source-active messages from looping. Router R accepts source-active messages originated by Router S only from neighbor Router N or an MSDP mesh group member.

1. S -----> N -----> R

Router R (the router that accepts or rejects active-source messages) locates its MSDP peer-RPF neighbor (Router N) deterministically. A series of rules is applied in a particular order to received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected.

The six rules applied to source-active messages originating at Router S received at Router R from Router N are as follows:

1. If Router N originated the source-active message (Router N is Router S), then Router N is also the peer-RPF neighbor, and its source-active messages are accepted.
2. If Router N is a member of the Router R mesh group, or is the configured peer, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
3. If Router N is the BGP next hop of the active multicast RPF route toward Router S (Router N installed the route on Router R), then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
4. If Router N is an external BGP (EBGP) or internal BGP (IBGP) peer of Router R, and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router N's AS number, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
5. If Router N uses the same next hop as the next hop to Router S, then Router N is the peer-RPF neighbor, and its source-active messages are accepted.
6. If Router N fits none of these criteria, then Router N is not an MSDP peer-RPF neighbor, and its source-active messages are rejected.

The MSDP peers that receive source-active TLVs can be constrained by BGP reachability information. If the AS path of the network layer reachability information (NLRI) contains the receiving peer's AS number prepended second to last, the sending peer is using the receiving peer as a next hop for this

source. If the split horizon information is not being received, the peer can be pruned from the source-active TLV distribution list.

For information about configuring MSDP mesh groups, see ["Example: Configuring MSDP with Active Source Limits and Mesh Groups" on page 559](#).

SEE ALSO

| [Configuring MSDP | 546](#)

Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the `msdp` statement:

```
msdp {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation (disable | enable);
  export [ policy-names ];
  group group-name {
    ... group-configuration ...
  }
  hold-time seconds;
  import [ policy-names ];
  local-address address;
  keep-alive seconds;
  peer address {
    ... peer-configuration ...
  }
  rib-group group-name;
  source ip-prefix</prefix-length> {
    active-source-limit {
      maximum number;
      threshold number;
    }
  }
  sa-hold-time seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
```

```

    flag flag <flag-modifier> <disable>;
}
group group-name {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    mode (mesh-group | standard);
    peer address {
        ... same statements as at the [edit protocols msdp peer address] hierarchy level
        shown just following ...
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
peer address {
    disable;
    active-source-limit {
        maximum number;
        threshold number;
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    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 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]

By default, MSDP is disabled.

SEE ALSO

| [Example: Configuring MSDP with Active Source Limits and Mesh Groups](#) | 559

Example: Configuring MSDP in a Routing Instance

IN THIS SECTION

- [Requirements](#) | 548
- [Overview](#) | 548
- [Configuration](#) | 551
- [Verification](#) | 557

This example shows how to configure MSDP in a VRF instance.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Enable PIM. See "[PIM Overview](#)" on page 275.

Overview

IN THIS SECTION

- [Topology](#) | 551

You can configure MSDP in the following types of instances:

- Forwarding
- No forwarding
- Virtual router
- VPLS
- VRF

The main use of MSDP in a routing instance is to support anycast RPs in the network, which allows you to configure redundant RPs. Anycast RP addressing requires MSDP support to synchronize the active sources between RPs.

This example includes the following MSDP settings.

- **authentication-key**—By default, multicast routers accept and process any properly formatted MSDP messages from the configured peer address. This default behavior might violate the security policies in many organizations because MSDP messages by definition come from another routing domain beyond the control of the security practices of the multicast router's organization.

The router can authenticate MSDP messages using the TCP message digest 5 (MD5) signature option for MSDP peering sessions. This authentication provides protection against spoofed packets being introduced into an MSDP peering session. Two organizations implementing MSDP authentication must decide on a human-readable key on both peers. This key is included in the MD5 signature computation for each MSDP segment sent between the two peers.

You configure an MSDP authentication key on a per-peer basis, whether the MSDP peer is defined in a group or individually. If you configure different authentication keys for the same peer one in a group and one individually, the individual key is used.

The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (" ").

Adding, removing, or changing an MSDP authentication key in a peering session resets the existing MSDP session and establishes a new session between the affected MSDP peers. This immediate session termination prevents excessive retransmissions and eventual session timeouts due to mismatched keys.

- **import** and **export**—All routing protocols use the routing table to store the routes that they learn and to determine which routes they advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in, and retrieve from, the routing table.

You can configure routing policy globally, for a group, or for an individual peer. This example shows how to configure the policy for an individual peer.

If you configure routing policy at the group level, each peer in a group inherits the group's routing policy.

The **import** statement applies policies to source-active messages being imported into the source-active cache from MSDP. The **export** statement applies policies to source-active messages being exported from the source-active cache into MSDP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found for the import policy, MSDP shares with the routing table only those routes that were learned from MSDP routers. If no match is found for the export policy, the default MSDP export policy is applied to entries in the source-active cache. See [Table 16 on page 550](#) for a list of match conditions.

Table 16: MSDP Source-Active Message Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address
neighbor	Neighbor address (the source address in the IP header of the source-active message)
route-filter	Multicast group address embedded in the source-active message
source-address-filter	Multicast source address embedded in the source-active message

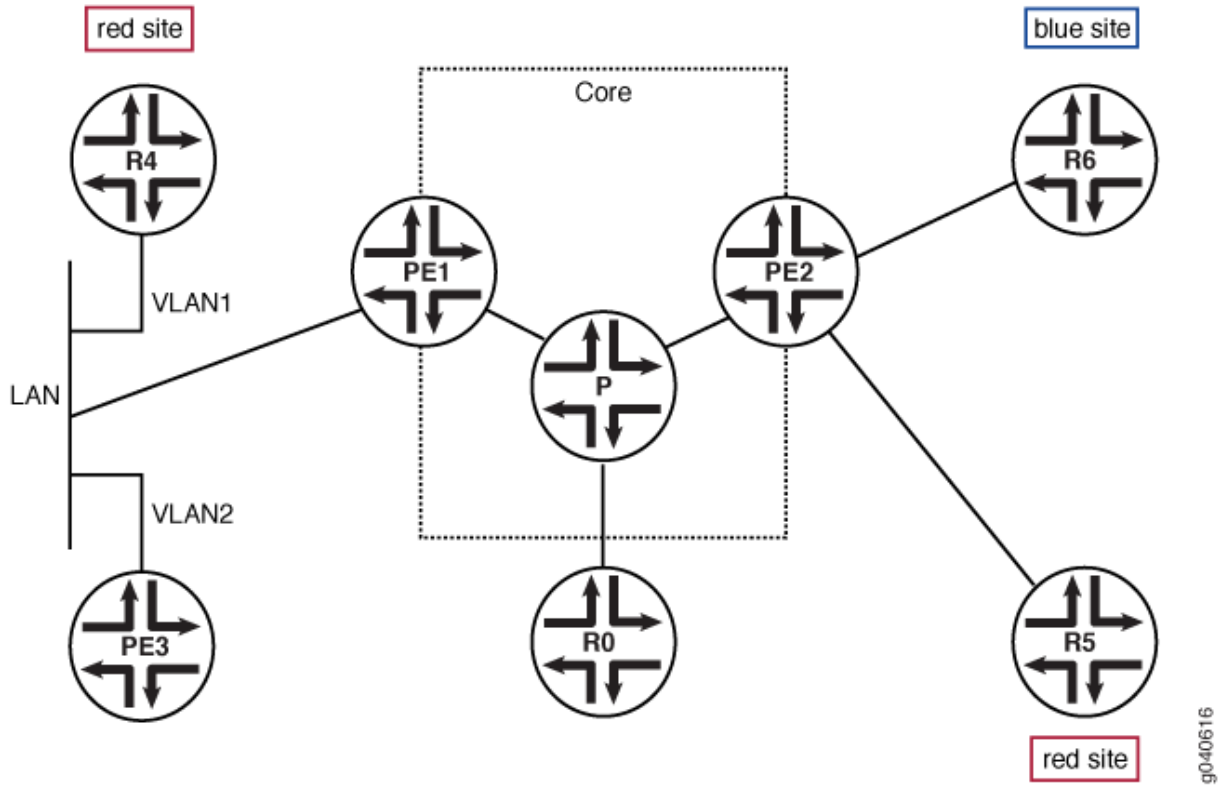
- **local-address**—Identifies the address of the router you are configuring as an MSDP router (the local router). When you configure MSDP, the **local-address** statement is required. The router must also be a Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP).
- **peer**—An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. You must configure at least one peer for MSDP to function. When you configure MSDP, the **peer** statement is required. The router must also be a Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP).

You can arrange MSDP peers into groups. Each group must contain at least one peer. Arranging peers into groups is useful if you want to block sources from some peers and accept them from others, or set tracing options on one group and not others. This example shows how to configure the MSDP peers in groups. If you configure MSDP peers in a group, each peer in a group inherits all group-level options.

Topology

Figure 77 on page 551 shows the topology for this example.

Figure 77: MSDP in a VRF Instance Topology



Configuration

IN THIS SECTION

- Procedure | 552
- Results | 555

Procedure

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.

```

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 sa-filter term bad-groups from route-filter 224.0.1.2/32
exact
set policy-options policy-statement sa-filter term bad-groups from route-filter 224.77.0.0/16
orlonger
set policy-options policy-statement sa-filter term bad-groups then reject
set policy-options policy-statement sa-filter term bad-sources from source-address-filter
10.0.0.0/8 orlonger
set policy-options policy-statement sa-filter term bad-sources from source-address-filter
127.0.0.0/8 orlonger
set policy-options policy-statement sa-filter term bad-sources then reject
set policy-options policy-statement sa-filter term accept-everything-else then accept
set routing-instances VPN-100 instance-type vrf
set routing-instances VPN-100 interface ge-0/0/0.100
set routing-instances VPN-100 interface lo0.100
set routing-instances VPN-100 route-distinguisher 10.255.120.36:100
set routing-instances VPN-100 vrf-target target:100:1
set routing-instances VPN-100 protocols ospf export bgp-to-ospf
set routing-instances VPN-100 protocols ospf area 0.0.0.0 interface lo0.100
set routing-instances VPN-100 protocols ospf area 0.0.0.0 interface ge-0/0/0.100
set routing-instances VPN-100 protocols pim rp static address 11.11.47.100
set routing-instances VPN-100 protocols pim interface lo0.100 mode sparse-dense
set routing-instances VPN-100 protocols pim interface lo0.100 version 2
set routing-instances VPN-100 protocols pim interface ge-0/0/0.100 mode sparse-dense
set routing-instances VPN-100 protocols pim interface ge-0/0/0.100 version 2
set routing-instances VPN-100 protocols msdp export sa-filter
set routing-instances VPN-100 protocols msdp import sa-filter
set routing-instances VPN-100 protocols msdp group 100 local-address 10.10.47.100
set routing-instances VPN-100 protocols msdp group 100 peer 10.255.120.39 authentication-key
"New York"
set routing-instances VPN-100 protocols msdp group to_pe local-address 10.10.47.100
set routing-instances VPN-100 protocols msdp group to_pe peer 11.11.47.100

```

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 [Junos OS CLI User Guide](#).

To configure an MSDP routing instance:

1. Configure the BGP export policy.

```
[edit policy-options]
user@host# set policy-statement bgp-to-ospf term 1 from protocol bgp
user@host# set policy-statement bgp-to-ospf term 1 then accept
```

2. Configure a policy that filters out certain source and group addresses and accepts all other source and group addresses.

```
[edit policy-options]
user@host# set policy-statement sa-filter term bad-groups from route-filter 224.0.1.2/32
exact
user@host# set policy-statement sa-filter term bad-groups from route-filter 224.0.1.2/32
exact
user@host# set policy-statement sa-filter term bad-groups from route-filter 224.77.0.0/16
orlonger
user@host# set policy-statement sa-filter term bad-groups then reject
user@host# set policy-statement sa-filter term bad-sources from source-address-filter
10.0.0.0/8 orlonger
user@host# set policy-statement sa-filter term bad-sources from source-address-filter
127.0.0.0/8 orlonger
user@host# set policy-statement sa-filter term bad-sources then reject
user@host# set policy-statement sa-filter term accept-everything-else then accept
```

3. Configure the routing instance type and interfaces.

```
[edit routing-instances]
user@host# set VPN-100 instance-type vrf
user@host# set VPN-100 interface ge-0/0/0.100
user@host# set VPN-100 interface lo0.100
```

4. Configure the routing instance route distinguisher and VRF target.

```
[edit routing-instances]
user@host# set VPN-100 route-distinguisher 10.255.120.36:100
user@host# set VPN-100 vrf-target target:100:1
```

5. Configure OSPF in the routing instance.

```
[edit routing-instances]
user@host# set VPN-100 protocols ospf export bgp-to-ospf
user@host# set VPN-100 protocols ospf area 0.0.0.0 interface lo0.100
user@host# set VPN-100 protocols ospf area 0.0.0.0 interface ge-0/0/0.100
```

6. Configure PIM in the routing instance.

```
[edit routing-instances]
user@host# set VPN-100 protocols pim rp static address 11.11.47.100
user@host# set VPN-100 protocols pim interface lo0.100 mode sparse-dense
user@host# set VPN-100 protocols pim interface lo0.100 version 2
user@host# set VPN-100 protocols pim interface ge-0/0/0.100 mode sparse-dense
user@host# set VPN-100 protocols pim interface ge-0/0/0.100 version 2
```

7. Configure MSDP in the routing instance.

```
[edit routing-instances]
user@host# set VPN-100 protocols msdp export sa-filter
user@host# set VPN-100 protocols msdp import sa-filter
user@host# set VPN-100 protocols msdp group 100 local-address 10.10.47.100
user@host# set VPN-100 protocols msdp group 100 peer 10.255.120.39 authentication-key "New
York"
[edit routing-instances]
user@host# set VPN-100 protocols msdp group to_pe local-address 10.10.47.100
[edit routing-instances]
user@host# set VPN-100 protocols msdp group to_pe peer 11.11.47.100
```

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

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** command and the **show routing-instances** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
}
policy-statement sa-filter {
  term bad-groups {
    from {
      route-filter 224.0.1.2/32 exact;
      route-filter 224.77.0.0/16 orlonger;
    }
    then reject;
  }
  term bad-sources {
    from {
      source-address-filter 10.0.0.0/8 orlonger;
      source-address-filter 127.0.0.0/8 orlonger;
    }
    then reject;
  }
  term accept-everything-else {
    then accept;
  }
}
```

```
}  
}
```

```
user@host# show routing-instances
```

```
VPN-100 {  
  instance-type vrf;  
  interface ge-0/0/0.100; ## 'ge-0/0/0.100' is not defined  
  interface lo0.100; ## 'lo0.100' is not defined  
  route-distinguisher 10.255.120.36:100;  
  vrf-target target:100:1;  
  protocols {  
    ospf {  
      export bgp-to-ospf;  
      area 0.0.0.0 {  
        interface lo0.100;  
        interface ge-0/0/0.100;  
      }  
    }  
    pim {  
      rp {  
        static {  
          address 11.11.47.100;  
        }  
      }  
      interface lo0.100 {  
        mode sparse-dense;  
        version 2;  
      }  
      interface ge-0/0/0.100 {  
        mode sparse-dense;  
        version 2;  
      }  
    }  
    msdp {  
      export sa-filter;  
      import sa-filter;  
      group 100 {  
        local-address 10.10.47.100;  
        peer 10.255.120.39 {  
          authentication-key "Hashed key found - Replaced with $ABC123abc123"; ##
```

```
SECRET-DATA
```

```
    }
  }
  group to_pe {
    local-address 10.10.47.100;
    peer 11.11.47.100;
  }
}
}
```

Verification

To verify the configuration, run the following commands:

- `show msdp instance VPN-100`
- `show msdp source-active VPN-100`
- `show multicast usage instance VPN-100`
- `show route table VPN-100.inet.4`

SEE ALSO

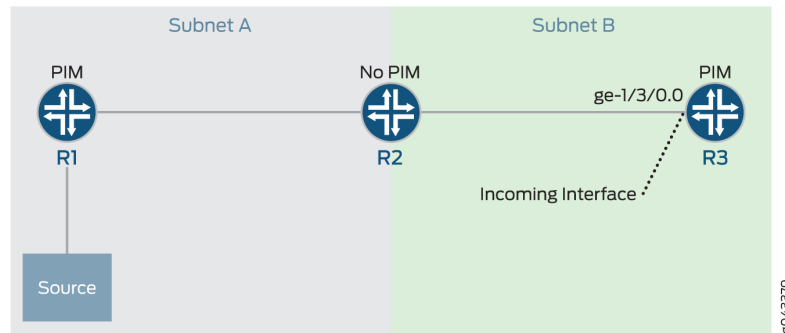
[Configuring Local PIM RPs | 343](#)

[Example: Configuring PIM Anycast With or Without MSDP | 358](#)

Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept multicast traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. [Figure 78 on page 558](#) shows such a topology, where R2 connects to the R1 source on one subnet, and to the incoming interface on R3 (ge-1/3/0.0 in the figure) on another subnet.

Figure 78: Accepting Multicast Traffic from a Remote Source



In this topology R2 is a pass-through device not running PIM, so R3 is the first hop router for multicast packets sent from R1. Because R1 and R3 are in different subnets, the default behavior of R3 is to disregard R1 as a remote source. You can have R3 accept multicast traffic from R1, however, by enabling `accept-remote-source` on the target interface.

To accept traffic from a remote source:

1. Identify the router and physical interface that you want to receive multicast traffic from the remote source.
2. Configure the interface to accept traffic from the remote source.

```
[edit protocols pim interface ge-1/3/0.0]
user@host# set accept-remote-source
```



NOTE: If the interface you identified is not the only path from the remote source, you need to ensure that it is the best path. For example you can configure a static route on the receiver side PE router to the source, or you can prepend the AS path on the other possible routes:

```
[edit policy-options policy-statement as-path-prepend term prepend]
user@host# set from route-filter 192.168.0.0/16 orlonger
user@host# set from route-filter 172.16.0.0/16 orlonger
user@host# set then as-path-prepend "1 1 1"
```

3. Commit the configuration changes.
4. Confirm that the interface you configured accepts traffic from the remote source.

```
user@host# show pim statistics
```


SEE ALSO

[Example: Allowing MBGP MVPN Remote Sources | 812](#)

Understanding Prepending AS Numbers to BGP AS Paths

show pim statistics

Example: Configuring MSDP with Active Source Limits and Mesh Groups

IN THIS SECTION

- [Requirements | 559](#)
- [Overview | 559](#)
- [Configuration | 563](#)
- [Verification | 566](#)

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Enable PIM sparse mode. See "[PIM Overview](#)" on page 275.
- Configure the router as a PIM sparse-mode RP. See "[Configuring Local PIM RPs](#)" on page 343.

Overview

IN THIS SECTION

- [Topology | 563](#)

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early detection (RED) to drop some but not all MSDP active source messages.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.

By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.



NOTE: The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is

configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.



CAUTION: When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.



NOTE: An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the **default-peer** statement.

[Table 17 on page 561](#) explains how flooding is handled by peers in this example. .

Table 17: Source-Active Message Flooding Explanation

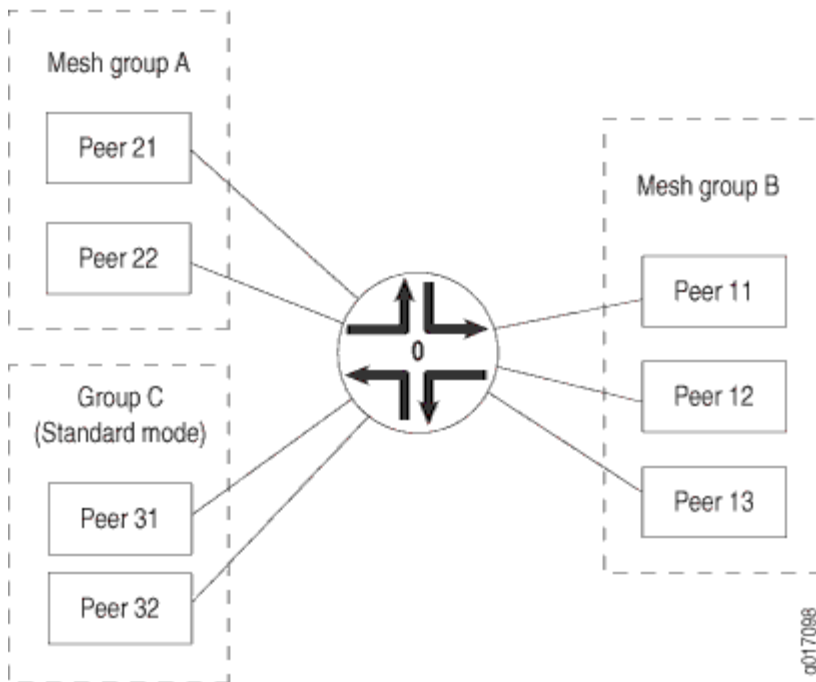
Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 21	Peer 11, Peer 12, Peer 13, Peer 31, Peer 32	Peer 22
Peer 11	Peer 21, Peer 22, Peer 31, Peer 32	Peer 12, Peer 13

Table 17: Source-Active Message Flooding Explanation (Continued)

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 31	Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32	-

Figure 79 on page 562 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

Figure 79: Source-Active Message Flooding



This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables.

As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the **MSDP-group** group are mesh group members.
- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group **MSDP-group**.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

Topology

Configuration

IN THIS SECTION

- Procedure | 564
- Results | 565

Procedure

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.

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500
```

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 [Junos OS CLI User Guide](#).

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

```
[edit protocols msdp]
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500
```

3. (Optional) Configure the threshold at which warning messages are logged and the amount of time between log messages.

```
[edit protocols msdp]
user@host# set active-source-limit log-warning 80
user@host# set active-source-limit log-interval 20
```

4. Configure the mesh group.

```
[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3
```

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

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** command.

```
user@host# show protocols
msdp {
  data-encapsulation disable;
  active-source-limit {
    maximum 10000;
  }
  peer 10.0.0.1 {
    active-source-limit {
      maximum 5000;
      threshold 4000;
    }
  }
  source 10.1.0.0/16 {
    active-source-limit {
      maximum 500;
    }
  }
}
```

```

}
group MSDP-group {
  mode mesh-group;
  local-address 10.1.2.3;
  peer 10.10.10.10 {
    active-source-limit {
      maximum 7500;
    }
  }
}
}
}

```

Verification

To verify the configuration, run the following commands:

- **show msdp source-active**
- **show msdp statistics**

SEE ALSO

[Filtering MSDP SA Messages | 379](#)

[Configuring Local PIM RPs | 343](#)

Tracing MSDP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
general	Trace general events.
keepalive	Trace keepalive messages.

(Continued)

Flag	Description
normal	Trace normal events.
packets	Trace all MSDP packets.
policy	Trace policy processing.
route	Trace MSDP changes to the routing table.
source-active	Trace source-active packets.
source-active-request	Trace source-active request packets.
source-active-response	Trace source-active response packets.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the `traceoptions` statement at the `[edit routing-options]` hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the MSDP trace file.

```
[edit protocols msdp group groupa traceoptions]  
user@host# set file msdp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols msdp group groupa traceoptions]  
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols msdp group groupa traceoptions]  
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols msdp group groupa traceoptions]  
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

```
[edit protocols msdp group groupa traceoptions]  
user@host# set flag source-active | match 230.0.0.3
```

7. View the trace file.

```
user@host> file list /var/log  
user@host> file show /var/log/msdp-trace
```

SEE ALSO

[Understanding MSDP | 544](#)

[Tracing and Logging Junos OS Operations](#)

[Junos OS Administration Library for Routing Devices](#)

Disabling MSDP

To disable MSDP on the router, include the disable statement:

```
disable;
```

You can disable MSDP globally for all peers, for all peers in a group, or for an individual peer.

- Globally for all MSDP peers at the following hierarchy levels:
 - [edit protocols msdp]
 - [edit logical-systems *logical-system-name* protocols msdp]
 - [edit routing-instances *routing-instance-name* protocols msdp]
 - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]
- For all peers in a group at the following hierarchy levels:
 - [edit protocols msdp group *group-name*]
 - [edit logical-systems *logical-system-name* protocols msdp group *group-name*]
 - [edit routing-instances *routing-instance-name* protocols msdp group *group-name*]
 - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*]
- For an individual peer at the following hierarchy levels:
 - [edit protocols msdp peer *address*]
 - [edit protocols msdp group *group-name* peer *address*]
 - [edit logical-systems *logical-system-name* protocols msdp peer *address*]
 - [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*]
 - [edit routing-instances *routing-instance-name* protocols msdp peer *address*]
 - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*]
 - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*]

If you disable MSDP at the group level, each peer in the group is disabled.

Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
  interface-routes {
    rib-group ifrg;
  }
  rib-groups {
    ifrg {
      import-rib [inet.0 inet.2];
    }
    mcrg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
}
protocols {
  bgp {
    group lab {
      type internal;
      family any;
      neighbor 192.168.6.18 {
        local-address 192.168.6.17;
      }
    }
  }
}
pim {
  dense-groups {
    224.0.1.39/32;
    224.0.1.40/32;
  }
  rib-group mcrg;
  rp {
    local {
      address 192.168.1.1;
    }
  }
  interface all {
    mode sparse-dense;
  }
}
```

```

        version 1;
    }
}
msdp {
    rib-group mcrg;
    group lab {
        peer 192.168.6.18 {
            local-address 192.168.6.17;
        }
    }
}
}
}

```

RELATED DOCUMENTATION

| [Understanding MSDP | 544](#)

Configuring Multiple Instances of MSDP

MSDP instances are supported for VRF instance types. For QFX5100, QFX5110, QFX5200, and EX9200 switches, MSDP instances are also supported for default and virtual router instance types. You can configure multiple instances of MSDP to support multicast over VPNs.

To configure multiple instances of MSDP, include the following statements:

```

routing-instances {
    routing-instance-name {
        interface interface-name;
        instance-type vrf;
        route-distinguisher (as-number: number | ip-address: number);
        vrf-import [ policy-names ];
        vrf-export [ policy-names ];
        protocols {
            msdp {
                ... msdp-configuration ...
            }
        }
    }
}

```

```
}  
}
```

You can include the statements at the following hierarchy levels:

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

RELATED DOCUMENTATION

[Examples: Configuring MSDP | 544](#)

[Junos OS MPLS Applications User Guide](#)

[Junos OS VPNs Library for Routing Devices](#)

Handling Session Announcements with SAP and SDP

IN THIS CHAPTER

- [Configuring the Session Announcement Protocol | 573](#)
- [Verifying SAP and SDP Addresses and Ports | 575](#)

Configuring the Session Announcement Protocol

IN THIS SECTION

- [Understanding SAP and SDP | 573](#)
- [Configuring the Session Announcement Protocol | 574](#)

Understanding SAP and SDP

Session announcements are handled by two protocols: the Session Announcement Protocol (SAP) and the Session Description Protocol (SDP). These two protocols display multicast session names and correlate the names with multicast traffic.

SDP is a session directory protocol that is used for multimedia sessions. It helps advertise multimedia conference sessions and communicates setup information to participants who want to join the session. SDP simply formats the session description. It does not incorporate a transport protocol. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.

SAP is a session directory announcement protocol that SDP uses as its transport protocol.

For information about supported standards for SAP and SDP, see ["Supported IP Multicast Protocol Standards" on page 21](#).

Configuring the Session Announcement Protocol

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.

The SAP and SDP protocols associate multicast session names with multicast traffic addresses. Only SAP has configuration parameters that users can change. Enabling SAP allows the router to receive announcements about multimedia and other multicast sessions.

Junos OS supports the following SAP and SDP standards:

- RFC 2327, *SDP Session Description Protocol*
- RFC 2974, *Session Announcement Protocol*

To enable SAP and the receipt of session announcements, include the `sap` statement:

```
sap {  
  disable;  
  listen address <port port>;  
}
```

You can include this statement at the following hierarchy levels:

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

By default, SAP listens to the address and port 224.2.127.254:9875 for session advertisements. To add other addresses or pairs of address and port, include one or more `listen` statements.

Sessions established by SDP, SAP's higher-layer protocol, time out after 60 minutes.

To monitor the operation, use the `show sap listen` command.

SEE ALSO

| `show sap listen`

Verifying SAP and SDP Addresses and Ports

IN THIS SECTION

- Purpose | 575
- Action | 575
- Meaning | 576

Purpose

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

Action

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

Sample Output

command-name

```
user@host> show sap listen
Group Address  Port
224.2.127.254 9875
```

Meaning

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

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

Facilitating Multicast Delivery Across Unicast-Only Networks with AMT

IN THIS CHAPTER

- [Example: Configuring Automatic IP Multicast Without Explicit Tunnels | 577](#)

Example: Configuring Automatic IP Multicast Without Explicit Tunnels

IN THIS SECTION

- [Understanding AMT | 577](#)
- [AMT Applications | 579](#)
- [AMT Operation | 580](#)
- [Configuring the AMT Protocol | 581](#)
- [Configuring Default IGMP Parameters for AMT Interfaces | 584](#)
- [Example: Configuring the AMT Protocol | 587](#)

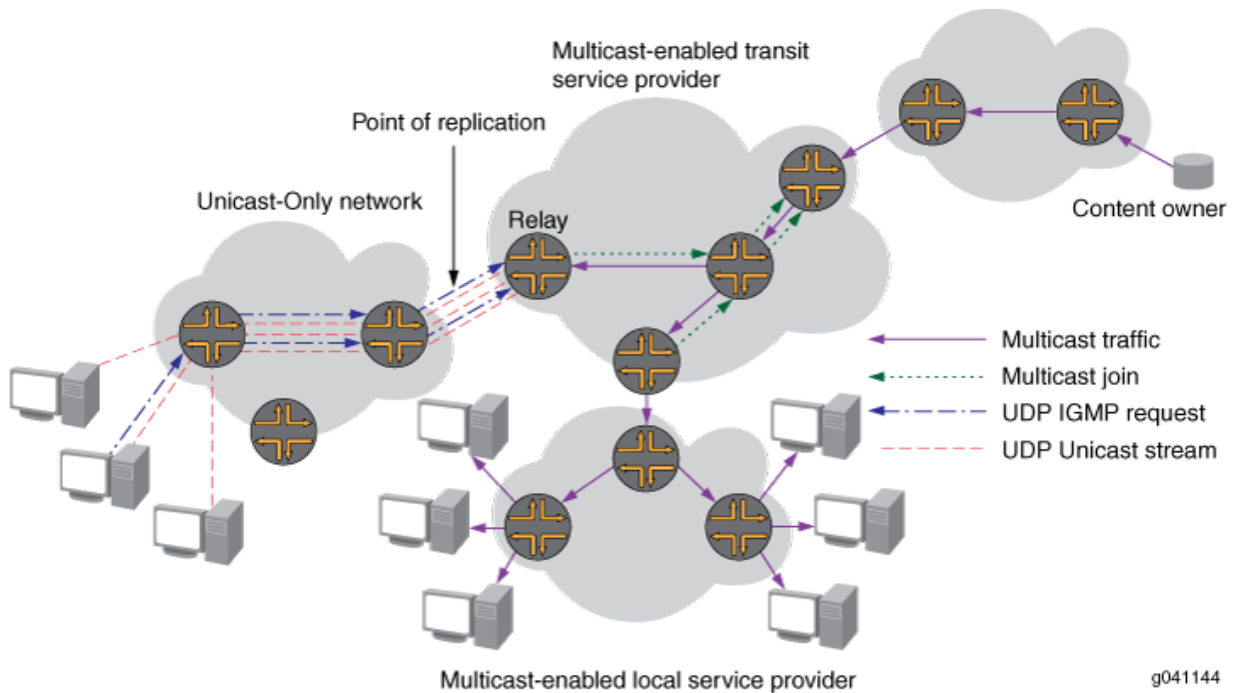
Understanding AMT

Automatic Multicast Tunneling (AMT) facilitates dynamic multicast connectivity between multicast-enabled networks across islands of unicast-only networks. Such connectivity enables service providers, content providers, and their customers to participate in delivering multicast traffic even if they lack end-to-end multicast connectivity.

AMT is supported on MX Series Ethernet Services Routers with Modular Port Concentrators (MPCs) that are running Junos 13.2 or later. AMT is also supported on i-chip based MPCs. AMT supports graceful restart (GR) but does not support *graceful Routing Engine switchover* (GRES).

AMT dynamically establishes unicast-encapsulated tunnels between well-known multicast-enabled relay points (AMT relays) and network points reachable only through unicast (AMT gateways). [Figure 80 on page 578](#) shows the Automatic Multicast Tunneling Connectivity.

Figure 80: Automatic Multicast Tunneling Connectivity



The AMT protocol provides discovery and handshaking between relays and gateways to establish tunnels dynamically without requiring explicit per-tunnel configuration.

AMT relays are typically routers with native IP multicast connectivity that aggregate a potentially large number of AMT tunnels.

The Junos OS implementation supports the following AMT relay functions:

- IPv4 multicast traffic and IPv4 encapsulation
- Well-known sources located on the multicast network
- Prevention of denial-of-service attacks by quickly discarding multicast packets that are sourced through a gateway.
- Per-route replication to the full fan-out of all AMT tunnels desired
- The ability to collect normal interface statistics on AMT tunnels

Multicast sources located behind AMT gateways are not supported.

AMT supports PIM sparse mode. AMT does not support dense mode operation.

AMT Applications

Transit service providers have a challenge in the Internet because many local service providers are not multicast-enabled. The challenge is how to entice content owners to transmit video and other multicast traffic across their backbones. The cost model for the content owners might be prohibitively high if they have to pay for unicast streams for the majority of their subscribers.

Until more local providers are multicast-enabled, there is a transition strategy proposed by the Internet Engineering Task Force (IETF) and implemented in open source software. This strategy is called Automatic IP Multicast Without Explicit Tunnels (AMT). AMT involves setting up relays at peering points in multicast networks that can be reached from gateways installed on hosts connected to unicast networks.

Without AMT, when a user who is connected to a unicast-only network wants to receive multicast content, the content owner can allow the user to join through unicast. However, the content owner incurs an added cost because the owner needs extra bandwidth to support the unicast subscribers.

AMT allows any host to receive multicast. On the client end is an AMT gateway that is a single host. Once the gateway has located an AMT relay, which might be a host but is more typically a router, the gateway periodically sends Internet Group Management Protocol (IGMP) messages over a dynamically created UDP tunnel to the relay. AMT relays and gateways cooperate to transmit multicast traffic sourced within the multicast network to end-user sites. AMT relays receive the traffic natively and unicast-encapsulate it to gateways. This allows anyone on the Internet to create a dynamic tunnel to download multicast data streams.

With AMT, a multicast-enabled service provider can offer multicast services to a content owner. When a customer of the unicast-only local provider wants to receive the content and subscribes using an AMT join, the multicast-enabled transit provider can then efficiently transport the content to the unicast-only local provider, which sends it on to the end user.

AMT is an excellent way for transit service providers (who can get access to the content, but do not have many end users) to provide multicast service to content owners, where it would not otherwise be economically feasible. It is also a useful transition strategy for local service providers who do not yet have multicast support on all downstream equipment.

AMT is also useful for connecting two multicast-enabled service providers that are separated by a unicast-only service provider.

Similarly, AMT can be used by local service providers whose networks are multicast-enabled to tunnel multicast traffic over legacy edge devices such as digital subscriber line access multiplexers (DSLAMs) that have limited multicast capabilities.

Technical details of the implementation of AMT are as follows:

- A three-way handshake is used to join groups from unicast receivers to prevent spoofing and denial-of-service (DoS) attacks.
- An AMT relay acting as a replication server joins the multicast group and translates multicast traffic into multiple unicast streams.
- The discovery mechanism uses anycast, enabling the discovery of the relay that is closest to the gateway in the network topology.
- An AMT gateway acting as a client is a host that joins the multicast group.
- Tunnel count limits on relays can limit bandwidth usage and avoid degradation of service.

AMT is described in detail in Internet draft draft-ietf-mboned-auto-multicast-10.txt, *Automatic IP Multicast Without Explicit Tunnels (AMT)*.

AMT Operation

AMT is used to create multicast tunnels dynamically between multicast-enabled networks across islands of unicast-only networks. To do this, several steps occur sequentially.

1. The AMT relay (typically a router) advertises an anycast address prefix and route into the unicast routing infrastructure.
2. The AMT gateway (a host) sends AMT relay discovery messages to the nearest AMT relay reachable across the unicast-only infrastructure. To reduce the possibility of replay attacks or dictionary attacks, the relay discovery messages contain a cryptographic nonce. A cryptographic nonce is a random number used only once.
3. The closest relay in the topology receives the AMT relay discovery message and returns the nonce from the discovery message in an AMT relay advertisement message. This enables the gateway to learn the relay's unique IP address. The AMT relay now has an address to use for all subsequent (S,G), entries it will join.
4. The AMT gateway sends an AMT request message to the AMT relay's unique IP address to begin the process of joining the (S,G).
5. The AMT relay sends an AMT membership query back to the gateway.
6. The AMT gateway receives the AMT query message and sends an AMT membership update message containing the IGMP join messages.
7. The AMT relay sends a join message toward the source to build a native multicast tree in the native multicast infrastructure.
8. As packets are received from the source, the AMT relay replicates the packets to all interfaces in the outgoing interface list, including the AMT tunnel. The multicast traffic is then encapsulated in unicast AMT multicast data messages.

9. To maintain state in the AMT relay, the AMT gateway sends periodic AMT membership updates.
10. After the tunnel is established, the AMT tunnel state is refreshed with each membership update message sent. The timeout for the refresh messages is 240 seconds.
11. When the AMT gateway leaves the group, the AMT relay can free resources associated with the tunnel.

Note the following operational details:

- The AMT relay creates an AMT pseudo interface (tunnel interface). AMT tunnel interfaces are implemented as generic UDP encapsulation (**ud**) logical interfaces. These logical interfaces have the identifier format **ud-fpc/pic/port.unit**.
- All multicast packets (data and control) are encapsulated in unicast packets. UDP encapsulation is used for all AMT control and data packets using the IANA reserved UDP port number (2268) for AMT.
- The AMT relay maintains a receiver list for each multicast session. The relay maintains the multicast state for each gateway that has joined a particular group or (S,G) pair.

Configuring the AMT Protocol

To configure the AMT protocol, include the `amt` statement:

```
amt {
  relay {
    accounting;
    family {
      inet {
        anycast-prefix ip-prefix</prefix-length>;
        local-address ip-address;
      }
    }
    secret-key-timeout minutes;
    tunnel-limit number;
  }
  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 protocols]
- [edit logical-systems *logical-system-name* protocols]
- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



NOTE: In the following example, only the [edit protocols] hierarchy is identified. The minimum configuration to enable AMT is to specify the AMT local address and the AMT anycast prefix.

1. To enable the MX Series router to create the UDP encapsulation (**ud**) logical interfaces, include the bandwidth statement and specify the bandwidth in gigabits per second.

```
[edit chassis fpc 0 pic 1]
user@host# set tunnel-services bandwidth 1g
```

2. Specify the local address by including the local-address statement at the [edit protocols amt relay family inet] hierarchy level.

```
[edit protocols amt relay family inet]
user@host# set local-address 192.168.7.1
```

The local address is used as the IP source of AMT control messages and the source of AMT data tunnel encapsulation. The local address can be configured on any active interface. Typically, the IP address of the router's **lo0.0** loopback interface is used for configuring the AMT local address in the default routing instance, and the IP address of the router's **lo0.n** loopback interface is used for configuring the AMT local address in VPN routing instances.

3. Specify the AMT anycast address by including the anycast-prefix statement at the [edit protocols amt relay family inet] hierarchy level.

```
[edit protocols amt relay family inet]
user@host# set anycast-prefix 192.168.0.0/16
```

The AMT anycast prefix is advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways. Typically, the router's **lo0.0** interface loopback address is used for configuring the AMT anycast prefix in the default routing instance, and the router's **lo0.n** loopback address is used for configuring the AMT anycast prefix in VPN routing instances. However, the anycast address can be either the primary or secondary **lo0.0** loopback address.

Ensure that your unicast routing protocol advertises the AMT anycast prefix in the route advertisements. If the AMT anycast prefix is advertised by BGP, ensure that the local autonomous system (AS) number for the AMT relay router is in the AS path leading to the AMT anycast prefix.

4. (Optional) Enable AMT accounting.

```
[edit protocols amt relay]
user@host# set accounting
```

5. (Optional) Specify the AMT secret key timeout by including the `secret-key-timeout` statement at the `[edit protocols amt relay]` hierarchy level. In the following example, the secret key timeout is configured to be 120 minutes.

```
[edit protocols amt relay]
user@host# set secret-key-timeout 120
```

The secret key is used to generate the AMT Message Authentication Code (MAC). Setting the secret key timeout shorter might improve security, but it consumes more CPU resources. The default is 60 minutes.

6. (Optional) Specify an AMT tunnel device by including the `tunnel-devices` statement at the `[edit protocols amt relay]` hierarchy level.

```
[edit protocols amt relay]
user@host# set tunnel-device 1
```

7. (Optional) Specify an AMT tunnel limit by including the `tunnel-limit` statement at the `[edit protocols amt relay]` hierarchy level. In the following example, the AMT tunnel limit is 12.

```
[edit protocols amt relay]
user@host# set tunnel-limit 12
```

The tunnel limit configures the static upper limit to the number of AMT tunnels that can be established. When the limit is reached, new AMT relay discovery messages are ignored.

8. Trace AMT protocol traffic by specifying options to the `traceoptions` statement at the `[edit protocols amt]` hierarchy level. Options applied at the AMT protocol level trace only AMT traffic. In the following example, all AMT packets are logged to the file `amt-log`.

```
[edit protocols amt]
user@host# set traceoptions file amt-log
user@host# set traceoptions flag packets
```



NOTE: For AMT operation, configure the PIM rendezvous point address as the primary loopback address of the AMT relay.

SEE ALSO

[CLI Explorer](#)

Configuring Default IGMP Parameters for AMT Interfaces

You can optionally configure default IGMP parameters for all AMT tunnel interfaces. Although, typically you do not need to change the values. To configure default IGMP attributes of all AMT relay tunnels, include the `amt` statement:

```
amt {
  relay {
    defaults {
      (accounting | no-accounting);
      group-policy [ policy-names ];
      query-interval seconds;
      query-response-interval seconds;
      robust-count number;
      ssm-map ssm-map-name;
      version version;
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols igmp]`

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

The IGMP statements included at the [edit protocols igmp amt relay defaults] hierarchy level have the same syntax and purpose as IGMP statements included at the [edit protocols igmp] or [edit protocols igmp interface *interface-name*] hierarchy levels. These statements are as follows:

- You can collect IGMP join and leave event statistics. To enable the collection of IGMP join and leave event statistics for all AMT interfaces, include the accounting statement:

```
user@host# set protocols igmp amt relay defaults accounting
```

- After enabling IGMP accounting, you must configure the router to filter the recorded information to a file or display it to a terminal. You can archive the events file.
- To disable the collection of IGMP join and leave event statistics for all AMT interfaces, include the no-accounting statement:

```
user@host# set protocols igmp amt relay defaults no-accounting
```

- You can filter unwanted IGMP reports at the interface level. To filter unwanted IGMP reports, define a policy to match only IGMP group addresses (for IGMPv2) by using the policy's route-filter statement to match the group address. Define the policy to match IGMP (S,G) addresses (for IGMPv3) by using the policy's route-filter statement to match the group address and the policy's source-address-filter statement to match the source address. In the following example, the **amt_reject** policy is created to match both the group and source addresses.

```
user@host# set policy-options policy-statement amt_reject from route-filter 224.1.1.1/32 exact
user@host# set policy-options policy-statement amt_reject from source-address-filter
192.168.0.0/16 orlonger
user@host# set policy-options policy-statement amt_reject then reject
```

- To apply the IGMP report filtering on the interface where you prefer not to receive specific group or (S,G) reports, include the group-policy statement. The following example applies the **amt_reject** policy to all AMT interfaces.

```
user@host# set protocols igmp amt relay defaults group-policy amt_reject
```

- You can change the IGMP query interval for all AMT interfaces to reduce or increase the number of host query messages sent. In AMT, host query messages are sent in response to membership request messages from the gateway. The query interval configured on the relay must be compatible with the membership request timer configured on the gateway. To modify this interval, include the `query-interval` statement. The following example sets the host query interval to 250 seconds.

```
user@host# set protocols igmp amt relay defaults query-interval 250
```

The IGMP querier router periodically sends general host-query messages. These messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

- You can change the IGMP query response interval. The query response interval multiplied by the robust count is the maximum amount of time that can elapse between the sending of a host query message by the querier router and the receipt of a response from a host. Varying this interval allows you to adjust the number of IGMP messages on the AMT interfaces. To modify this interval, include the `query-response-interval` statement. The following example configures the query response interval to 20 seconds.

```
user@host# set protocols igmp amt relay defaults query-response-interval 20
```

- You can change the IGMP robust count. The robust count is used to adjust for the expected packet loss on the AMT interfaces. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork. To modify the robust count, include the `robust-count` statement. The following example configures the robust count to 3.

```
user@host# set protocols igmp amt relay defaults robust-count 3
```

The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3.

- On a shared network running IGMPv2, when the query router receives an IGMP leave message, it must send an IGMP group query message for a specified number of times. The number of IGMP group query messages sent is determined by the robust count. The interval between query messages is determined by the last member query interval. Also, the IGMPv2 query response interval is multiplied by the robust count to determine the maximum amount of time between the sending of a host query message and receipt of a response from a host.

For more information about the IGMPv2 robust count, see RFC 2236, *Internet Group Management Protocol, Version 2*.

- In IGMPv3 a change of interface state causes the system to immediately transmit a state-change report from that interface. If the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3

the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

For more information about the IGMPv3 robust count, see RFC 3376, *Internet Group Management Protocol, Version 3*.

- You can apply a source-specific multicast (SSM) map to an AMT interface. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report, which allows hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4).

In this example, you create a policy to match the 232.1.1.1/32 group address for translation to IGMPv3. Then you define the SSM map that associates the policy with the 192.168.43.66 source address where these group addresses are found. Finally, you apply the SSM map to all AMT interfaces.

```
user@host# set policy-options policy-statement ssm-policy-example term A from route-filter
232.1.1.1/32 exact
user@host# set policy-options policy-statement ssm-policy-example term A then accept
user@host# set routing-options multicast ssm-map ssm-map-example policy ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-example source 192.168.43.66
user@host# set protocols igmp amt relay defaults ssm-map ssm-map-example
```

SEE ALSO

Specifying Log File Size, Number, and Archiving Properties

[Junos OS Administration Library for Routing Devices](#)

Example: Configuring the AMT Protocol

IN THIS SECTION

- [Requirements | 588](#)
- [Overview | 588](#)
- [Configuration | 589](#)
- [Verification | 592](#)

This example shows how to configure the Automatic Multicast Tunneling (AMT) Protocol to facilitate dynamic multicast connectivity between multicast-enabled networks across islands of unicast-only networks.

Requirements

Before you begin:

- Configure the router interfaces.
- 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](#)" on [page 27](#) and "[Understanding MLD](#)" on [page 59](#).

Overview

IN THIS SECTION

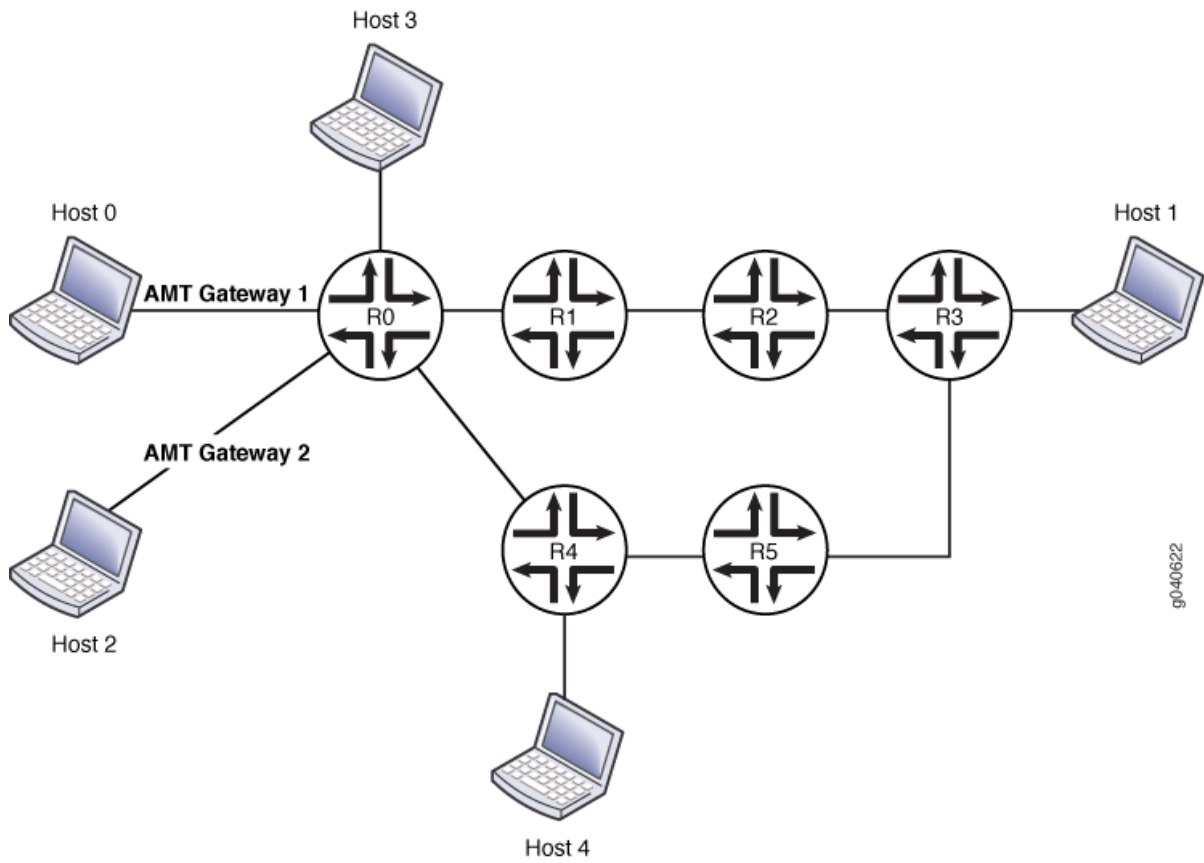
- [Topology | 588](#)

In this example, Host 0 and Host 2 are multicast receivers in a unicast cloud. Their default gateway devices are AMT gateways. R0 and R4 are configured with unicast protocols only. R1, R2, R3, and R5 are configured with PIM multicast. Host 1 is a source in a multicast cloud. R0 and R5 are configured to perform AMT relay. Host 3 and Host 4 are multicast receivers (or sources that are directly connected to receivers). This example shows R1 configured with an AMT relay local address and an anycast prefix as its own loopback address. The example also shows R0 configured with tunnel services enabled.

Topology

[Figure 81 on page 589](#) shows the topology used in this example.

Figure 81: AMT Gateway Topology



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 590](#)
- [Procedure | 590](#)
- [Results | 591](#)

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.

```
set protocols amt traceoptions file amt.log
set protocols amt traceoptions flag errors
set protocols amt traceoptions flag packets detail
set protocols amt traceoptions flag route detail
set protocols amt traceoptions flag state detail
set protocols amt traceoptions flag tunnels detail
set protocols amt relay family inet anycast-prefix 10.10.10.10/32
set protocols amt relay family inet local-address 10.255.112.201
set protocols amt relay tunnel-limit 10
set protocols pim interface all mode sparse-dense
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
```

Procedure

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 [Junos OS CLI User Guide](#).

To configure the AMT protocol on R1:

1. Configure AMT tracing operations.

```
[edit protocols amt traceoptions]
user@host# set file amt.log
user@host# set flag errors
user@host# set flag packets detail
user@host# set flag route detail
user@host# set flag state detail
user@host# set flag tunnels detail
```


2. Configure the AMT relay settings.

```
[edit protocols amt relay]
user@host# set relay family inet anycast-prefix 10.10.10.10/32
user@host# set family inet local-address 10.255.112.201
user@host# set tunnel-limit 10
```

3. Configure PIM on R1's interfaces.

```
[edit protocols pim]
set interface all mode sparse-dense
set interface all version 2
set interface fxp0.0 disable
```

4. Enable tunnel functionality.

```
[edit chassis]
set fpc 0 pic 0 tunnel-services bandwidth 1g
```

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

```
user@host# commit
```

Results

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

```
user@host# show chassis
fpc 0 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
```

```
}  
}
```

```
user@host# show protocols  
amt {  
  traceoptions {  
    file amt.log;  
    flag errors;  
    flag packets detail;  
    flag route detail;  
    flag state detail;  
    flag tunnels detail;  
  }  
  relay {  
    family {  
      inet {  
        anycast-prefix 10.10.10.10/32;  
        local-address 10.255.112.201;  
      }  
    }  
    tunnel-limit 10;  
  }  
}  
pim {  
  interface all {  
    mode sparse-dense;  
    version 2;  
  }  
  interface fxp0.0 {  
    disable;  
  }  
}
```

Verification

To verify the configuration, run the following commands:

- **show amt statistics**
- **show amt summary**
- **show amt tunnel**

RELATED DOCUMENTATION

| [Understanding AMT | 577](#)

Routing Content to Densely Clustered Receivers with DVMRP

IN THIS CHAPTER

- [Examples: Configuring DVMRP | 594](#)

Examples: Configuring DVMRP

IN THIS SECTION

- [Understanding DVMRP | 594](#)
- [Configuring DVMRP | 595](#)
- [Example: Configuring DVMRP | 596](#)
- [Example: Configuring DVMRP to Announce Unicast Routes | 600](#)
- [Tracing DVMRP Protocol Traffic | 606](#)

Understanding DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

The Distance Vector Multicast Routing Protocol (DVMRP) is a distance-vector routing protocol that provides connectionless datagram delivery to a group of hosts across an internetwork. DVMRP is a distributed protocol that dynamically generates IP multicast delivery trees by using a technique called reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces. These mechanisms allow the formation of shortest-path trees, which are used to reach all group members from each network source of multicast traffic.

DVMRP is designed to be used as an interior gateway protocol (IGP) within a multicast domain.

Because not all IP routers support native multicast routing, DVMRP includes direct support for tunneling IP multicast datagrams through routers. The IP multicast datagrams are encapsulated in unicast IP packets and addressed to the routers that do support native multicast routing. DVMRP treats tunnel interfaces and physical network interfaces the same way.

DVMRP routers dynamically discover their neighbors by sending neighbor probe messages periodically to an IP multicast group address that is reserved for all DVMRP routers.

Configuring DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Distance Vector Multicast Routing Protocol (DVMRP) is the first of the multicast routing protocols and has a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G).

To configure the Distance Vector Multicast Routing Protocol (DVMRP), include the `dvmrp` statement:

```
dvmrp {
  disable;
  export [ policy-names ];
  import [ policy-names ];
  interface interface-name {
    disable;
    hold-time seconds;
    metric metric;
    mode (forwarding | unicast-routing);
  }
  rib-group group-name;
  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 protocols]

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

By default, DVMRP is disabled.

Example: Configuring DVMRP

IN THIS SECTION

- Requirements | 596
- Overview | 596
- Configuration | 598
- Verification | 600

This example shows how to use DVMRP to announce routes used for multicast routing as well as multicast data forwarding.

Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview

DVMRP is a distance vector protocol for multicast. It is similar to RIP, in that both RIP and DVMRP have issues with scalability and robustness. PIM domains are more commonly used than DVMRP domains. In some environments, you might need to configure interoperability with DVMRP.

This example includes the following DVMRP settings:

- **protocols dvmrp rib-group**—Associates the **dvmrp-rib** routing table group with the DVMRP protocol to enable multicast RPF lookup.

- **protocols dvmrp interface**—Configures the DVMRP interface. The interface of a DVMRP router can be either a physical interface to a directly attached subnetwork or a tunnel interface to another multicast-capable area of the Multicast Backbone (*MBone*). The DVMRP hold-time period is the amount of time that a neighbor is to consider the sending router (this router) to be operative (up). The default hold-time period is 35 seconds.
- **protocols dvmrp interface hold-time**—The DVMRP hold-time period is the amount of time that a neighbor is to consider the sending router (this router) to be operative (up). The default hold-time period is 35 seconds.
- **protocols dvmrp interface metric**—All interfaces can be configured with a metric specifying cost for receiving packets on a given interface. The default metric is 1.

For each source network reported, a route metric is associated with the unicast route being reported. The metric is the sum of the interface metrics between the router originating the report and the source network. A metric of 32 marks the source network as unreachable, thus limiting the breadth of the DVMRP network and placing an upper bound on the DVMRP convergence time.

- **routing-options rib-groups**—Enables DVMRP to access route information from the unicast routing table, **inet.0**, and from a separate routing table that is reserved for DVMRP. In this example, the first routing table group named **ifrg** contains local interface routes. This ensures that local interface routes get added to both the **inet.0** table for use by unicast protocols and the **inet.2** table for multicast RPF check. The second routing table group named **dvmrp-rib** contains **inet.2** routes.

DVMRP needs to access route information from the unicast routing table, **inet.0**, and from a separate routing table that is reserved for DVMRP. You need to create the routing table for DVMRP and to create groups of routing tables so that the routing protocol process imports and exports routes properly. We recommend that you use routing table **inet.2** for DVMRP routing information.

- **routing-options interface-routes**— After defining the **ifrg** routing table group, use the `interface-routes` statement to insert interface routes into the **ifrg** group—in other words, into both **inet.0** and **inet.2**. By default, interface routes are imported into routing table **inet.0** only.
- **sap**—Enables the Session Directory Announcement Protocol (SAP) and the Session Directory Protocol (SDP). Enabling SAP allows the router to receive announcements about multimedia and other multicast sessions.

SAP always listens to the address and port 224.2.127.254:9875 for session advertisements. To add other addresses or pairs of address and port, include one or more `listen` statements.

Sessions learned by SDP, SAP's higher-layer protocol, time out after 60 minutes.

Configuration

IN THIS SECTION

- Procedure | 598
- Results | 599

Procedure

CLI Quick Configuration

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

```
set routing-options interface-routes rib-group inet ifrg
set routing-options rib-groups ifrg import-rib inet.0
set routing-options rib-groups ifrg import-rib inet.2
set routing-options rib-groups dvmrp-rib export-rib inet.2
set routing-options rib-groups dvmrp-rib import-rib inet.2
set protocols sap
set protocols dvmrp rib-group dvmrp-rib
set protocols dvmrp interface ip-0/0/0.0 metric 5
set protocols dvmrp interface ip-0/0/0.0 hold-time 40
```

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 [Junos OS CLI User Guide](#).

To configure an MSDP routing instance:

1. Create the routing tables for DVMRP routes.

```
[edit routing-options]
user@host# set interface-routes rib-group inet ifrg
user@host# set rib-groups ifrg import-rib [ inet.0 inet.2 ]
```



```
user@host# set rib-groups dvmrp-rib import-rib inet.2
user@host# set rib-groups dvmrp-rib export-rib inet.2
```

2. Configure SAP and SDP.

```
[edit protocols]
user@host# set sap
```

3. Enable DVMRP on the router and associate the **dvmrp-rib** routing table group with DVMRP to enable multicast RPF checks.

```
[edit protocols]
user@host# set dvmrp rib-group dvmrp-rib
```

4. Configure the DVMRP interface with a hold-time value and a metric. This example shows an IP-over-IP encapsulation tunnel interface.

```
[edit protocols]
user@host# set dvmrp interface ip-0/0/0.0
user@host# set dvmrp interface ip-0/0/0.0 hold-time 40
user@host# set dvmrp interface ip-0/0/0.0 metric 5
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-options` command and the `show protocols` command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
interface-routes {
  rib-group inet ifrg;
}
rib-groups {
```

```
ifrg {
    import-rib [ inet.0 inet.2 ];
}
dvmp-rib {
    export-rib inet.2;
    import-rib inet.2;
}
}
```

```
user@host# show protocols
sap;
dvmp {
    rib-group dvmp-rib;
    interface ip-0/0/0.0 {
        metric 5;
        hold-time 40;
    }
}
```

Verification

To verify the configuration, run the following commands:

- **show dvmp interfaces**
- **show dvmp neighbors**

Example: Configuring DVMP to Announce Unicast Routes

IN THIS SECTION

- [Requirements | 601](#)
- [Overview | 601](#)
- [Configuration | 602](#)
- [Verification | 605](#)

Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

This example shows how to use DVMRP to announce unicast routes used solely for multicast reverse-path forwarding (RPF) to set up the multicast control plane.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview

IN THIS SECTION

- [Topology | 602](#)

DVMRP has two modes. Forwarding mode is the default mode. In forwarding mode, DVMRP is responsible for the multicast control plane and multicast data forwarding. In the nondefault mode (which is shown in this example), DVMRP does not forward multicast data traffic. This mode is called unicast routing mode because in this mode DVMRP is only responsible for announcing unicast routes used for multicast RPF—in other words, for establishing the control plane. To forward multicast data, enable Protocol Independent Multicast (PIM) on the interface. If you have configured PIM on the interface, as shown in this example, you can configure DVMRP in unicast-routing mode only. You cannot configure PIM and DVMRP in forwarding mode at the same time.

This example includes the following settings:

- **policy-statement dvmrp-export**—Accepts static default routes.
- **protocols dvmrp export dvmrp-export**—Associates the **dvmrp-export** policy with the DVMRP protocol.

All routing protocols use the routing table to store the routes that they learn and to determine which routes they advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table. Import and export policies are always

from the point of view of the routing table. So the **dvmrp-export** policy exports static default routes from the routing table and accepts them into DVMRP.

- **protocols dvmrp interface all mode unicast-routing**—Enables all interfaces to announce unicast routes used solely for multicast RPF.
- **protocols dvmrp rib-group inet dvmrp-rg**—Associates the **dvmrp-rib** routing table group with the DVMRP protocol to enable multicast RPF checks.
- **protocols pim rib-group inet pim-rg**—Associates the **pim-rg** routing table group with the PIM protocol to enable multicast RPF checks.
- **routing-options rib inet.2 static route 0.0.0.0/0 discard**—Redistributes static routes to all DVMRP neighbors. The **inet.2** routing table stores unicast IPv4 routes for multicast RPF lookup. The **discard** statement silently drops packets without notice.
- **routing-options rib-groups dvmrp-rg import-rib inet.2**—Creates the routing table for DVMRP to ensure that the routing protocol process imports routes properly.
- **routing-options rib-groups dvmrp-rg export-rib inet.2**—Creates the routing table for DVMRP to ensure that the routing protocol process exports routes properly.
- **routing-options rib-groups pim-rg import-rib inet.2**—Enables access to route information from the routing table that stores unicast IPv4 routes for multicast RPF lookup. In this example, the first routing table group named **pim-rg** contains local interface routes. This ensures that local interface routes get added to the **inet.2** table.

Topology

Configuration

IN THIS SECTION

- Procedure | 603
- Results | 604

Procedure

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.

```
set policy-options policy-statement dvmrp-export term 10 from protocol static
set policy-options policy-statement dvmrp-export term 10 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement dvmrp-export term 10 then accept
set protocols dvmrp rib-group inet
set protocols dvmrp rib-group dvmrp-rg
set protocols dvmrp export dvmrp-export
set protocols dvmrp interface all mode unicast-routing
set protocols dvmrp interface fxp0.0 disable
set protocols pim rib-group inet pim-rg
set protocols pim interface all
set routing-options rib inet.2 static route 0.0.0.0/0 discard
set routing-options rib-groups pim-rg import-rib inet.2
set routing-options rib-groups dvmrp-rg export-rib inet.2
set routing-options rib-groups dvmrp-rg import-rib inet.2
```

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 [Junos OS CLI User Guide](#).

To configure an MSDP routing instance:

1. Configure the routing options.

```
[edit routing-options]
[edit routing -options]
user@host# set rib inet.2 static route 0.0.0.0/0 discard
user@host# set rib-groups pim-rg import-rib inet.2
user@host# set rib-groups dvmrp-rg import-rib inet.2
user@host# set rib-groups dvmrp-rg export-rib inet.2
```

2. Configure DVMRP.

```
[edit protocols]
user@host# set dvmrp rib-group inet dvmrp-rg
user@host# set dvmrp export dvmrp-export
user@host# set dvmrp interface all mode unicast-routing
user@host# set dvmrp interface fxp0 disable
```

3. Configure PIM so that PIM performs multicast data forwarding.

```
[edit protocols]
user@host# set pim rib-group inet pim-rg
user@host# set pim interface all
```

4. Configure the DVMRP routing policy.

```
[edit policy-options policy-statement dvmrp-export term 10]
user@host# set from protocol static
user@host# set from route-filter 0.0.0.0/0 exact
user@host# set then accept
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** command, the **show protocols** command, and the **show routing-options** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement dvmrp-export {
  term 10 {
    from {
      protocol static;
      route-filter 0.0.0.0/0 exact;
    }
  }
}
```

```
        then accept;
    }
}
```

```
user@host# show protocols
dvmrp {
    rib-group inet dvmrp-rg;
    export dvmrp-export;
    interface all {
        mode unicast-routing;
    }
    interface fxp0.0 {
        disable;
    }
}
pim {
    rib-group inet pim-rg;
    interface all;
}
```

```
user@host# show routing-options
rib inet.2 {
    static {
        route 0.0.0.0/0 discard;
    }
}
rib-groups {
    pim-rg {
        import-rib inet.2;
    }
    dvmrp-rg {
        export-rib inet.2;
        import-rib inet.2;
    }
}
```

Verification

To verify the configuration, run the following commands:

- `show dvmrp interfaces`
- `show pim statistics`

Tracing DVMRP Protocol Traffic

Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
general	Trace general flow.
graft	Trace graft messages.
neighbor	Trace neighbor probe packets.
normal	Trace normal events.
packets	Trace all DVMRP packets.
poison	Trace poison-route-reverse packets.
policy	Trace policy processing.
probe	Trace probe packets.
prune	Trace prune messages.
report	Trace membership report messages.

(Continued)

Flag	Description
route	Trace routing information.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on DVMRP packets of a particular type. To configure tracing operations for DVMRP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the DVMRP trace file.

```
[edit protocols dvmrp traceoptions]
user@host# set file dvmrp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols dvmrp traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols dvmrp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols dvmrp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular DVMRP neighbor. The following example shows how to trace neighbor probe packets that match the neighbor's IP address.

```
[edit protocols dvmrp traceoptions]
user@host# set flag neighbor | match 192.168.1.1
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/dvmrp-trace
```

SEE ALSO

[Tracing and Logging Junos OS Operations](#)

[Junos OS Administration Library for Routing Devices](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
16.1	Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

5

PART

Configuring Multicast VPNs

[Configuring Draft-Rosen Multicast VPNs | 610](#)

[Configuring Next-Generation Multicast VPNs | 715](#)

[Configuring PIM Join Load Balancing | 1059](#)

Configuring Draft-Rosen Multicast VPNs

IN THIS CHAPTER

- Draft-Rosen Multicast VPNs Overview | 610
- Example: Configuring Any-Source Draft-Rosen 6 Multicast VPNs | 611
- Example: Configuring a Specific Tunnel for IPv4 Multicast VPN Traffic (Using Draft-Rosen MVPNs) | 630
- Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs | 648
- Understanding Data MDTs | 661
- Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode | 663
- Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode | 669
- Examples: Configuring Data MDTs | 682

Draft-Rosen Multicast VPNs Overview

The Junos OS provides two types of draft-rosen multicast VPNs:

- Draft-rosen multicast VPNs with service provider tunnels operating in any-source multicast (ASM) mode (also referred to as *rosen 6* Layer 3 VPN multicast)—Described in RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and based on Section 2 of the IETF Internet draft **draft-rosen-vpn-mcast-06.txt**, *Multicast in MPLS/BGP VPNs* (expired April 2004).
- Draft-rosen multicast VPNs with service provider tunnels operating in source-specific multicast (SSM) mode (also referred to as *rosen 7* Layer 3 VPN multicast)—Described in RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and based on the IETF Internet draft **draft-rosen-vpn-mcast-07.txt**, *Multicast in MPLS/BGP IP VPNs*. Draft-rosen multicast VPNs with service provider tunnels operating in SSM mode do not require that the provider (P) routers maintain any VPN-specific Protocol-Independent Multicast (PIM) information.



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.

In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, the VPN is multicast-enabled and configured to use the Protocol Independent Multicast (PIM) protocol within the VPN and within the service provider (SP) network. A multicast-enabled VPN routing and forwarding (VRF) instance corresponds to a multicast domain (MD), and a PE router attached to a particular VRF instance is said to belong to the corresponding MD. For each MD there is a *default multicast distribution tree (MDT)* through the SP backbone, which connects all of the PE routers belonging to that MD. Any PE router configured with a default MDT group address can be the multicast source of one default MDT.

Draft-rosen MVPNs with service provider tunnels start by sending all multicast traffic over a default MDT, as described in section 2 of the IETF Internet draft **draft-rosen-vpn-mcast-06.txt** and section 7 of the IETF Internet draft **draft-rosen-vpn-mcast-07.txt**. This default mapping results in the delivery of packets to each provider edge (PE) router attached to the provider router even if the PE router has no receivers for the multicast group in that VPN. Each PE router processes the encapsulated VPN traffic even if the multicast packets are then discarded.

RELATED DOCUMENTATION

[Junos OS VPNs Library for Routing Devices](#)

Example: Configuring Any-Source Draft-Rosen 6 Multicast VPNs

IN THIS SECTION

- [Understanding Any-Source Multicast | 611](#)
- [Example: Configuring Any-Source Multicast for Draft-Rosen VPNs | 612](#)
- [Load Balancing Multicast Tunnel Interfaces Among Available PICs | 625](#)

Understanding Any-Source Multicast

Any-source multicast (ASM) is the form of multicast in which you can have multiple senders on the same group, as opposed to source-specific multicast where a single particular source is specified. The original multicast specification, RFC 1112, supports both the ASM many-to-many model and the SSM one-to-many model. For ASM, the (S,G) source, group pair is instead specified as (*,G), meaning that the multicast group traffic can be provided by multiple sources.

An ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of *source discovery* is a required function of the network itself.

In an environment where many sources come and go, such as for a video conferencing service, ASM is appropriate. Multicast source discovery appears to be an easy process, but in sparse mode it is not.

In PIM dense mode, it is simple enough to flood traffic to every router in the network so that every router learns the source address of the content for that multicast group.

However, in PIM sparse mode, traffic flooding is not a viable option as it presents issues related to scalability and network resource use.

SEE ALSO

[Example: Configuring Source-Specific Multicast Groups with Any-Source Override | 458](#)

[Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode | 663](#)

Example: Configuring Any-Source Multicast for Draft-Rosen VPNs

IN THIS SECTION

- [Requirements | 612](#)
- [Overview | 613](#)
- [Configuration | 616](#)
- [Verification | 625](#)

This example shows how to configure an any-source multicast VPN (MVPN) using dual PIM configuration with a customer RP and provider RP and mapping the multicast routes from customer to provider (known as *draft-rosen*). The Junos OS complies with RFC 4364 and Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*.

Requirements

Before you begin:

- Configure the router interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).

- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure the VPN. See the [Junos OS VPNs Library for Routing Devices](#).
- Configure the VPN import and VPN export policies. See [Configuring Policies for the VRF Table on PE Routers in VPNs](#) in the [Junos OS VPNs Library for Routing Devices](#).
- Make sure that the routing devices support multicast tunnel (**mt**) interfaces for encapsulating and de-encapsulating data packets into tunnels. See ["Tunnel Services PICs and Multicast" on page 315](#) and ["Load Balancing Multicast Tunnel Interfaces Among Available PICs" on page 611](#).

For multicast to work on draft-rosen Layer 3 VPNs, each of the following routers must have tunnel interfaces:

- Each provider edge (PE) router.
- Any provider (P) router acting as the RP.
- Any customer edge (CE) router that is acting as a source's DR or as an RP. A receiver's designated router does not need a Tunnel Services PIC.

Overview

IN THIS SECTION

- [Topology | 615](#)

Draft-rosen multicast virtual private networks (MVPNs) can be configured to support service provider tunnels operating in any-source multicast (ASM) mode or source-specific multicast (SSM) mode.

In this example, the term *multicast Layer 3 VPNs* is used to refer to draft-rosen MVPNs.

This example includes the following settings.

- **interface lo0.1**—Configures an additional unit on the loopback interface of the PE router. For the **lo0.1** interface, assign an address from the VPN address space. Add the **lo0.1** interface to the following places in the configuration:
 - VRF routing instance
 - PIM in the VRF routing instance
 - IGP and BGP policies to advertise the interface in the VPN address space


In multicast Layer 3 VPNs, the multicast PE routers must use the primary loopback address (or router ID) for sessions with their internal BGP peers. If the PE routers use a route reflector and the next hop is configured as **self**, Layer 3 multicast over VPN will not work, because PIM cannot transmit upstream interface information for multicast sources behind remote PEs into the network core. Multicast Layer 3 VPNs require that the BGP next-hop address of the VPN route match the BGP next-hop address of the loopback VRF instance address.

- **protocols pim interface**—Configures the interfaces between each provider router and the PE routers. On all CE routers, include this statement on the interfaces facing toward the provider router acting as the RP.
- **protocols pim mode sparse**—Enables PIM sparse mode on the **lo0** interface of all PE routers. You can either configure that specific interface or configure all interfaces with the `interface all` statement. On CE routers, you can configure sparse mode or sparse-dense mode.
- **protocols pim rp local**—On all routers acting as the RP, configure the address of the local **lo0** interface. The P router acts as the RP router in this example.
- **protocols pim rp static**—On all PE and CE routers, configure the address of the router acting as the RP.

It is possible for a PE router to be configured as the VPN customer RP (C-RP) router. A PE router can also act as the DR. This type of PE configuration can simplify configuration of customer DRs and VPN C-RPs for multicast VPNs. This example does not discuss the use of the PE as the VPN C-RP.

Figure 82 on page 614 shows multicast connectivity on the customer edge. In the figure, CE2 is the RP router. However, the RP router can be anywhere in the customer network.

Figure 82: Multicast Connectivity on the CE Routers

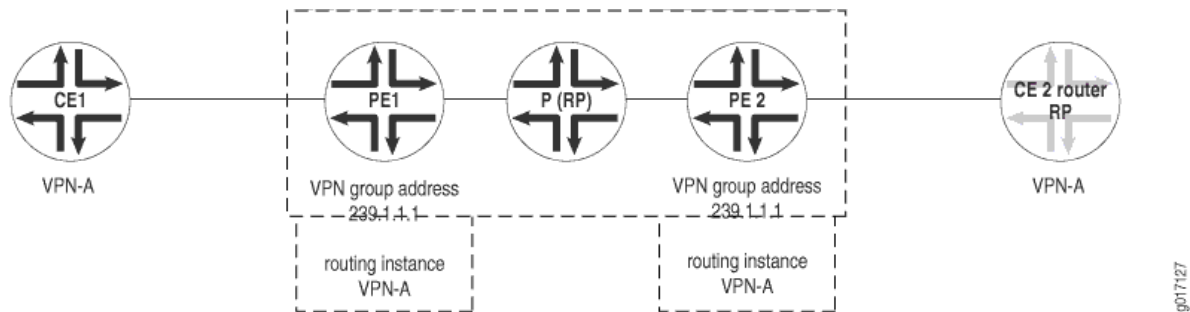


- **protocols pim version 2**—Enables PIM version 2 on the **lo0** interface of all PE routers and CE routers. You can either configure that specific interface or configure all interfaces with the `interface all` statement.
- **group-address**—In a routing instance, configure multicast connectivity for the VPN on the PE routers. Configure a VPN group address on the interfaces facing toward the router acting as the RP.

The PIM configuration in the VPN routing and forwarding (VRF) instance on the PE routers needs to match the master PIM instance on the CE router. Therefore, the PE router contains both a master PIM instance (to communicate with the provider core) and the VRF instance (to communicate with the CE routers).

VRF instances that are part of the same VPN share the same VPN group address. For example, all PE routers containing multicast-enabled routing instance VPN-A share the same VPN group address configuration. In [Figure 83 on page 615](#), the shared VPN group address configuration is 239.1.1.1.

Figure 83: Multicast Connectivity for the VPN




- **routing-instances *instance-name* protocols pim rib-group**—Adds the routing group to the VPN's VRF instance.
- **routing-options rib-groups**—Configures the multicast routing group.

Topology

This example describes how to configure multicast in PIM sparse mode for a range of multicast addresses for VPN-A as shown in [Figure 84 on page 615](#).

Figure 84: Customer Edge and Service Provider Networks



Configuration

IN THIS SECTION

- [Procedure | 616](#)
- [Results | 622](#)

Procedure

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 interfaces lo0 unit 0 family inet address 192.168.27.13/32 primary
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 1 family inet address 10.10.47.101/32
set protocols pim rp static address 10.255.71.47
set protocols pim interface fxp0.0 disable
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface t1-1/0/0:0.0
set routing-instances VPN-A interface lo0.1
set routing-instances VPN-A route-distinguisher 10.255.71.46:100
set routing-instances VPN-A vrf-import VPNA-import
set routing-instances VPN-A vrf-export VPNA-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 t1-1/0/0:0.0
set routing-instances VPN-A protocols ospf area 0.0.0.0 interface lo0.1
set routing-instances VPN-A protocols pim rib-group inet VPNA-mcast-rib
set routing-instances VPN-A protocols pim rp static address 10.255.245.91
set routing-instances VPN-A protocols pim interface t1-1/0/0:0.0 mode sparse
set routing-instances VPN-A protocols pim interface t1-1/0/0:0.0 version 2
set routing-instances VPN-A protocols pim interface lo0.1 mode sparse
set routing-instances VPN-A protocols pim interface lo0.1 version 2

```

```

set routing-instances VPN-A provider-tunnel pim-asm group-address 239.1.1.1
set routing-instances VPN-A protocols pim mvpn
set routing-options interface-routes rib-group inet VPNA-mcast-rib
set routing-options rib-groups VPNA-mcast-rib export-rib VPN-A.inet.2
set routing-options rib-groups VPNA-mcast-rib import-rib VPN-A.inet.2

```

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 [Junos OS CLI User Guide](#).

To configure multicast for draft-rosen VPNs:

1. Configure PIM on the P router.

```

[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set dense-groups 224.0.1.39/32
[edit protocols pim]
user@host# set dense-groups 224.0.1.40/32
[edit protocols pim]
user@host# set rp local address 10.255.71.47
[edit protocols pim]
user@host# set interface all mode sparse
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable

```

2. Configure PIM on the PE1 and PE2 routers. Specify a static RP—the P router (10.255.71.47).

```

[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.71.47
[edit protocols pim]
user@host# set interface interface all mode sparse
[edit protocols pim]
user@host# set interface interface all version 2

```

```
[edit protocols pim]
user@host# set interface fxp0.0 disable
[edit protocols pim]
user@host# exit
```

3. Configure PIM on CE1. Specify the RP address for the VPN RP—Router CE2 (10.255.245.91).

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.245.91
[edit protocols pim]
user@host# set interface all mode sparse
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
[edit protocols pim]
user@host# exit
```

4. Configure PIM on CE2, which acts as the VPN RP. Specify CE2's address (10.255.245.91).

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp local address 10.255.245.91
[edit protocols pim]
user@host# set interface all mode sparse
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
[edit protocols pim]
user@host# exit
```

5. On PE1, configure the routing instance (VPN-A) for the Layer 3 VPN.

```
[edit]
user@host# edit routing-instances VPN-A
[edit routing-instances VPN-A]
```

```

user@host# set instance-type vrf
[edit routing-instances VPN-A]
user@host# set interface t1-1/0/0:0.0
[edit routing-instances VPN-A]
user@host# set interface lo0.1
[edit routing-instances VPN-A]
user@host# set route-distinguisher 10.255.71.46:100
[edit routing-instances VPN-A]
user@host# set vrf-import VPNA-import
[edit routing-instances VPN-A]
user@host# set vrf-export VPNA-export

```

6. On PE1, configure the IGP policy to advertise the interfaces in the VPN address space.

```

[edit routing-instances VPN-A]
user@host# set protocols ospf export bgp-to-ospf
[edit routing-instances VPN-A]
user@host# set protocols ospf area 0.0.0.0 interface t1-1/0/0:0.0
[edit routing-instances VPN-A]
user@host# set protocols ospf area 0.0.0.0 interface lo0.1

```

7. On PE1, set the RP configuration for the VRF instance. The RP configuration within the VRF instance provides explicit knowledge of the RP address, so that the (*,G) state can be forwarded.

```

[edit routing-instances VPN-A]
user@host# set protocols pim mvpn
[edit routing-instances VPN-A]
user@host# set protocols provider-tunnel pim-asm group-address 239.1.1.1
[edit routing-instances VPN-A]
user@host# set protocols pim rp static address 10.255.245.91
[edit routing-instances VPN-A]
user@host# set protocols pim interface t1-1/0/0:0.0 mode sparse
[edit routing-instances VPN-A]
user@host# set protocols pim interface t1-1/0/0:0.0 version 2
[edit routing-instances VPN-A]
user@host# set protocols pim interface lo0.1 mode sparse
[edit routing-instances VPN-A]
user@host# set protocols pim interface lo0.1 version 2
[edit routing-instances VPN-A]
user@host# exit

```

8. On PE1, configure the loopback interfaces.

```
[edit]
user@host# edit interface lo0
[edit interface lo0]
user@host# set unit 0 family inet address 192.168.27.13/32 primary
[edit interface lo0]
user@host# set unit 0 family inet address 127.0.0.1/32
[edit interface lo0]
user@host# set unit 1 family inet address 10.10.47.101/32
[edit interface lo0]
user@host# exit
```

9. As you did for the PE1 router, configure the PE2 router.

```
[edit]
user@host# edit routing-instances VPN-A
[edit routing-instances VPN-A]
user@host# set instance-type vrf
[edit routing-instances VPN-A]
user@host# set interface t1-2/0/0:0.0
[edit routing-instances VPN-A]
user@host# set interface lo0.1
[edit routing-instances VPN-A]
user@host# set route-distinguisher 10.255.71.51:100
[edit routing-instances VPN-A]
user@host# set vrf-import VPNA-import
[edit routing-instances VPN-A]
user@host# set vrf-export VPNA-export
[edit routing-instances VPN-A]
user@host# set protocols ospf export bgp-to-ospf
[edit routing-instances VPN-A]
user@host# set protocols ospf area 0.0.0.0 interface t1-2/0/0:0.0
[edit routing-instances VPN-A]
user@host# set protocols ospf area 0.0.0.0 interface lo0.1
[edit routing-instances VPN-A]
user@host# set protocols pim rp static address 10.255.245.91
[edit routing-instances VPN-A]
user@host# set protocols pim mvpn
[edit routing-instances VPN-A]
user@host# set protocols pim interface t1-2/0/0:0.0 mode sparse
```

```

[edit routing-instances VPN-A]
user@host# set protocols pim interface lo0.1 mode sparse
[edit routing-instances VPN-A]
user@host# set protocols pim interface lo0.1 version 2
[edit routing-instances VPN-A]
user@host# set provider-tunnel pim-asm group-address 239.1.1.1
user@host# exit
[edit]
user@host# edit interface lo0
[edit interface lo0]
user@host# set unit 0 family inet address 192.168.27.14/32 primary
[edit interface lo0]
user@host# set unit 0 family inet address 127.0.0.1/32
[edit interface lo0]
user@host# set unit 1 family inet address 10.10.47.102/32

```

10. When one of the PE routers is running Cisco Systems IOS software, you must configure the Juniper Networks PE router to support this multicast interoperability requirement. The Juniper Networks PE router must have the **lo0.0** interface in the master routing instance and the **lo0.1** interface assigned to the VPN routing instance. You must configure the **lo0.1** interface with the same IP address that the **lo0.0** interface uses for BGP peering in the provider core in the master routing instance.

Configure the same IP address on the **lo0.0** and **lo0.1** loopback interfaces of the Juniper Networks PE router at the [edit interfaces lo0] hierarchy level, and assign the address used for BGP peering in the provider core in the master routing instance. In this alternate example, unit 0 and unit 1 are configured for Cisco IOS interoperability.

```

[edit interface lo0]
user@host# set unit 0 family inet address 192.168.27.14/32 primary
[edit interface lo0]
user@host# set unit 0 family inet address 127.0.0.1/32
[edit interface lo0]
user@host# set unit 1 family inet address 192.168.27.14/32
[edit interface lo0]
user@host# exit

```

11. Configure the multicast routing table group. This group accesses **inet.2** when doing RPF checks. However, if you are using **inet.0** for multicast RPF checks, this step will prevent your multicast configuration from working.

```
[edit]
user@host# edit routing-options
[edit routing-options]
user@host# set interface-routes rib-group inet VPNA-mcast-rib
[edit routing-options]
user@host# set rib-groups VPNA-mcast-rib export-rib VPN-A.inet.2
[edit routing-options]
user@host# set rib-groups VPNA-mcast-rib import-rib VPN-A.inet.2
[edit routing-options]
user@host# exit
```

12. Activate the multicast routing table group in the VPN's VRF instance.

```
[edit]
user@host# edit routing-instances VPN-A
[edit routing-instances VPN-A]
user@host# set protocols pim rib-group inet VPNA-mcast-rib
```

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

```
[edit routing-instances VPN-A]
user@host# commit
```

Results

Confirm your configuration by entering the `show interfaces`, `show protocols`, `show routing-instances`, and `show routing-options` commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration. This output shows the configuration on PE1.

```
user@host# show interfaces
lo0 {
  unit 0 {
    family inet {
      address 192.168.27.13/32 {
```



```

        primary;
    }
    address 127.0.0.1/32;
}
}
unit 1 {
    family inet {
        address 10.10.47.101/32;
    }
}
}
}

```

user@host# **show protocols**

```

pim {
    rp {
        static {
            address 10.255.71.47;
        }
    }
    interface fxp0.0 {
        disable;
    }
    interface all {
        mode sparse;
        version 2;
    }
}
}

```

user@host# **show routing-instances**

```

VPN-A {
    instance-type vrf;
    interface t1-1/0/0:0.0;
    interface lo0.1;
    route-distinguisher 10.255.71.46:100;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    provider-tunnel {
        pim-asm {
            group-address 239.1.1.1;
        }
    }
}

```

```

}
protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.0 {
      interface t1-1/0/0:0.0;
      interface lo0.1;
    }
  }
  pim {
    mvpn;
    rib-group inet VPNA-mcast-rib;
    rp {
      static {
        address 10.255.245.91;
      }
    }
    interface t1-1/0/0:0.0 {
      mode sparse;
      version 2;
    }
    interface lo0.1 {
      mode sparse;
      version 2;
    }
  }
}
}

```

```

user@host# show routing-options
interface-routes {
  rib-group inet VPNA-mcast-rib;
}
rib-groups {
  VPNA-mcast-rib {
    export-rib VPN-A.inet.2;
    import-rib VPN-A.inet.2;
  }
}

```

Verification

To verify the configuration, run the following commands:

1. Display multicast tunnel information and the number of neighbors by using the `show pim interfaces instance instance-name` command from the PE1 or PE2 router. When issued from the PE1 router, the output display is:

```
user@host> show pim interfaces instance VPN-A
Instance: PIM.VPN-A
Name                Stat Mode      IP V State Count DR address
lo0.1               Up   Sparse      4 2 DR         0 10.10.47.101
mt-1/1/0.32769      Up   Sparse      4 2 DR         1
mt-1/1/0.1081346    Up   Sparse      4 2 DR         0
pe-1/1/0.32769      Up   Sparse      4 1 P2P        0
t1-2/1/0:0.0        Up   Sparse      4 2 P2P        1
```

You can also display all PE tunnel interfaces by using the `show pim join` command from the provider router acting as the RP.

2. Display multicast tunnel interface information, DR information, and the PIM neighbor status between VRF instances on the PE1 and PE2 routers by using the `show pim neighbors instance instance-name` command from either PE router. When issued from the PE1 router, the output is as follows:

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
Interface           IP V Mode      Option      Uptime Neighbor addr
mt-1/1/0.32769      4 2           HPL         01:40:46 10.10.47.102
t1-1/0/0:0.0        4 2           HPL         01:41:41 192.168.196.178
```

SEE ALSO

[Example: Configuring PIM RPF Selection | 1167](#)

Load Balancing Multicast Tunnel Interfaces Among Available PICs

When you configure multicast on draft-rosen Layer 3 VPNs, multicast tunnel interfaces are automatically generated to encapsulate and de-encapsulate control and data traffic.

To generate multicast tunnel interfaces, a routing device must have one or more of the following tunnel-capable PICs:

- Adaptive Services PIC
- Multiservices PIC or Multiservices DPC
- Tunnel Services PIC
- On MX Series routers, a PIC created with the `tunnel-services` statement at the `[edit chassis fpc slot-number pic number]` hierarchy level



NOTE: A *routing device* is a router or an EX Series switch that is functioning as a router.

If a routing device has multiple such PICs, it might be important in your implementation to load balance the tunnel interfaces across the available tunnel-capable PICs.

The multicast tunnel interface that is used for encapsulation, `mt-[xxxxx]`, is in the range from 32,768 through 49,151. The interface `mt-[yyyyy]`, used for de-encapsulation, is in the range from 1,081,344 through 1,107,827. PIM runs only on the encapsulation interface. The de-encapsulation interface populates downstream interface information. For the default MDT, an instance's de-encapsulation and encapsulation interfaces are always created on the same PIC.

For each VPN, the PE routers build a multicast distribution tree within the service provider core network. After the tree is created, each PE router encapsulates all multicast traffic (data and control messages) from the attached VPN and sends the encapsulated traffic to the VPN group address. Because all the PE routers are members of the outgoing interface list in the multicast distribution tree for the VPN group address, they all receive the encapsulated traffic. When the PE routers receive the encapsulated traffic, they de-encapsulate the messages and send the data and control messages to the CE routers.

If a routing device has multiple tunnel-capable PICs (for example, two Tunnel Services PICs), the routing device load balances the creation of tunnel interfaces among the available PICs. However, in some cases (for example, after a reboot), a single PIC might be selected for all of the tunnel interfaces. This causes one PIC to have a heavy load, while other available PICs are underutilized. To prevent this, you can manually configure load balancing. Thus, you can configure and distribute the load uniformly across the available PICs.

The definition of a balanced state is determined by you and by the requirements of your Layer 3 VPN implementation. You might want all of the instances to be evenly distributed across the available PICs or across a configured list of PICs. You might want all of the encapsulation interfaces from all of the instances to be evenly distributed across the available PICs or across a configured list of PICs. If the bandwidth of each tunnel encapsulation interface is considered, you might choose a different distribution. You can design your load-balancing configuration based on each instance or on each routing device.



NOTE: In a Layer 3 VPN, each of the following routing devices must have at least one tunnel-capable PIC:

- Each provider edge (PE) router.
- Any provider (P) router acting as the RP.
- Any customer edge (CE) router that is acting as a source's DR or as an RP. A receiver's designated router does not need a tunnel-capable PIC.

To configure load balancing:

1. On an M Series or T Series router or on an EX Series switch, install more than one tunnel-capable PIC. (In some implementations, only one PIC is required. Load balancing is based on the assumption that a routing device has more than one tunnel-capable PIC.)
2. On an MX Series router, configure more than one tunnel-capable PIC.

```
[edit chassis fpc 0]
user@host# set pic 0 tunnel-services bandwidth 10g
user@host# set pic 1 tunnel-services bandwidth 10g
```

3. Configure Layer 3 VPNs as described in Example: Configuring Any-Source Multicast for Draft-Rosen VPNs.

```
[edit routing-instances vpn1]
user@host# set provider-tunnel pim-asm group-address 234.1.1.1
user@host# set protocols pim rp static address 10.255.72.48
user@host# set protocols pim interface fe-1/0/0.0
user@host# set protocols pim interface lo0.1
user@host# set protocols pim mvpn
```

4. For each VPN, specify a PIC list.

```
[edit routing-instances vpn1 protocols pim]
user@host# set tunnel-devices [ mt-1/1/0 mt-1/2/0 mt-2/0/0 ]
```

The physical position of the PIC in the routing device determines the multicast tunnel interface name. For example, if you have an Adaptive Services PIC installed in FPC slot 0 and PIC slot 0, the corresponding multicast tunnel interface name is `mt-0/0/0`. The same is true for Tunnel Services PICs, Multiservices PICs, and Multiservices DPCs.

In the `tunnel-devices` statement, the order of the PIC list that you specify does not impact how the interfaces are allocated. An instance uses all of the listed PICs to create default encapsulation and de-encapsulation interfaces, and data MDT encapsulation interfaces. The instance uses a round-robin approach to distributing the tunnel interfaces (default and data MDT) across the PIC list (or across the available PICs, in the absence of a PIC list).

For the first tunnel, the round-robin algorithm starts with the lowest-numbered PIC. The second tunnel is created on the next-lowest-numbered PIC, and so on, round and round. The selection algorithm works routing device-wide. The round robin does not restart at the lowest-numbered PIC for each new instance. This applies to both the default and data MDT tunnel interfaces.

If one PIC in the list fails, new tunnel interfaces are created on the remaining PICs in the list using the round-robin algorithm. If all the PICs in the list go down, all tunnel interfaces are deleted and no new tunnel interfaces are created. If a PIC in the list comes up from the down state and the restored PIC is the only PIC that is up, the interfaces are reassigned to the restored PIC. If a PIC in the list comes up from the down state and other PICs are already up, an interface reassignment is not done. However, when a new tunnel interface needs to be created, the restored PIC is available for the selection process. If you include in the PIC list a PIC that is not installed on the routing device, the PIC is treated as if it is present but in the down state.

To balance the interfaces among the instances, you can assign one PIC to each instance. For example, if you have `vpn1-10` and you have three PICs—for example, `mt-1/1/0`, `mt-1/2/0`, `mt-2/0/0`—you can configure `vpn1-4` to only use `mt-1/1/0`, `vpn5-7` to use `mt-1/2/0`, and `vpn8-10` to use `mt-2/0/0`.

5. Commit the configuration.

```
user@host# commit
```

When you commit a new PIC list configuration, all the multicast tunnel interfaces for the routing instance are deleted and re-created using the new PIC list.

6. If you reboot the routing device, some PICs come up faster than others. The difference can be minutes. Therefore, when the tunnel interfaces are created, the known PIC list might not be the same as when the routing device is fully rebooted. This causes the tunnel interfaces to be created on some but not all available and configured PICs. To remedy this situation, you can manually rebalance the PIC load.

Check to determine if a load rebalance is necessary.

```
user@host#> show interfaces terse | match mt-
mt-1/1/0          up    up
mt-1/1/0.32768   up    up    inet
mt-1/1/0.1081344 up    up    inet
mt-1/2/0          up    up
mt-1/2/0.32769  up    up    inet
```

```
mt-1/2/0.32770 up up inet
mt-1/2/0.32771 up up inet
```

The output shows that `mt-1/1/0` has only one tunnel encapsulation interface, while `mt-1/2/0` has three tunnel encapsulation interfaces. In a case like this, you might decide to rebalance the interfaces. As stated previously, encapsulation interfaces are in the range from 32,768 through 49,151. In determining whether a rebalance is necessary, look at the encapsulation interfaces only, because the default MDT de-encapsulation interface always resides on the same PIC with the default MDT encapsulation interface.

7. (Optional) Rebalance the PIC load.

```
user@host#> request pim multicast-tunnel rebalance instance vpn1
```

This command re-creates and rebalances all tunnel interfaces for a specific instance.

```
user@host#> request pim multicast-tunnel rebalance
```

This command re-creates and rebalances all tunnel interfaces for all routing instances.

8. Verify that the PIC load is balanced.

```
user@host#> show interfaces terse | match mt-
mt-1/1/0 up up
mt-1/1/0.32770 up up inet
mt-1/1/0.32768 up up inet
mt-1/1/0.1081344 up up inet
mt-1/2/0 up up
mt-1/2/0.32769 up up inet
mt-1/2/0.32771 up up inet
```

The output shows that `mt-1/1/0` has two encapsulation interfaces, and `mt-1/2/0` also has two encapsulation interfaces.

SEE ALSO

request pim multicast-tunnel rebalance

[CLI Explorer](#)

RELATED DOCUMENTATION

[Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs | 648](#)

Example: Configuring a Specific Tunnel for IPv4 Multicast VPN Traffic (Using Draft-Rosen MVPNs)

IN THIS SECTION

- [Requirements | 630](#)
- [Overview | 631](#)
- [PE Router Configuration | 632](#)
- [CE Device Configuration | 641](#)
- [Verification | 644](#)

This example shows how to configure different provider tunnels to carry IPv4 customer traffic in a multicast VPN network.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks devices: Two PE routers and two CE devices.
- Junos OS Release 11.4 or later running on the PE routers.
- The PE routers can be M Series Multiservice Edge Routers, MX Series Ethernet Services Routers, or T Series Core Routers.
- The CE devices can be switches (such as EX Series Ethernet Switches), or they can be routers (such as M Series, MX Series, or T Series platforms).

Overview

IN THIS SECTION

- [Topology Diagram | 632](#)

A multicast tunnel is a mechanism to deliver control and data traffic across the provider core in a multicast VPN. Control and data packets are transmitted over the multicast distribution tree in the provider core. When a service provider carries both IPv4 and IPv6 traffic from a single customer, it is sometimes useful to separate the IPv4 and IPv6 traffic onto different multicast tunnels within the customer VRF routing instance. Putting customer IPv4 and IPv6 traffic on two different tunnels provides flexibility and control. For example, it helps the service provider to charge appropriately, to manage and measure traffic patterns, and to have an improved capability to make decisions when deploying new services.

A draft-rosen 7 multicast VPN control plane is configured in this example. The control plane is configured to use source-specific multicast (SSM) mode. The provider tunnel is used for the draft-rosen 7 control traffic and IPv4 customer traffic.

This example uses the following statements to configure the draft-rosen 7 control plane and specify IPv4 traffic to be carried in the provider tunnel:

- `provider-tunnel pim-ssm family inet group-address 232.1.1.1`
- `pim mvpn family inet autodiscovery inet-mdt`
- `pim mvpn family inet6 disable`
- `mvpn family inet autodiscovery-only intra-as inclusive`
- `family inet-mdt signaling`


Note the following limitations:

- Junos OS does not currently support IPv6 with draft-rosen 6 or draft-rosen 7.
- Junos OS does not support more than two provider tunnels in a routing instance. For example, you cannot configure an RSVP-TE provider tunnel plus two MVPN provider tunnels.
- In a routing instance, you cannot configure both an any-source multicast (ASM) tunnel and an SSM tunnel.

Topology Diagram

Figure 85 on page 632 shows the topology used in this example.

Figure 85: Different Provider Tunnels for IPv4 Multicast VPN Traffic



PE Router Configuration

IN THIS SECTION

- [CLI Quick Configuration | 632](#)
- [Router PE1 | 635](#)
- [Results | 637](#)

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 PE1

```
set interfaces so-0/0/3 unit 0 family inet address 10.111.10.1/30
set interfaces so-0/0/3 unit 0 family mpls
set interfaces fe-1/1/2 unit 0 family inet address 10.10.10.1/30
set interfaces lo0 unit 0 family inet address 10.255.182.133/32 primary
set interfaces lo0 unit 1 family inet address 10.10.47.100/32
set routing-options router-id 10.255.182.133
set routing-options route-distinguisher-id 10.255.182.133
set routing-options autonomous-system 100
set routing-instances VPN-A instance-type vrf
```

```

set routing-instances VPN-A interface fe-1/1/2.0
set routing-instances VPN-A interface lo0.1
set routing-instances VPN-A provider-tunnel pim-ssm family inet group-address 232.1.1.1
set routing-instances VPN-A provider-tunnel mdt threshold group 224.1.1.0/24 source
10.240.0.242/32 rate 10
set routing-instances VPN-A provider-tunnel mdt tunnel-limit 20
set routing-instances VPN-A provider-tunnel mdt group-range 232.1.1.3/32
set routing-instances VPN-A vrf-target target:100:10
set routing-instances VPN-A vrf-table-label
set routing-instances VPN-A protocols ospf area 0.0.0.0 interface all
set routing-instances VPN-A protocols ospf export bgp-to-ospf
set routing-instances VPN-A protocols pim mvpn family inet autodiscovery inet-mdt
set routing-instances VPN-A protocols pim mvpn family inet6 disable
set routing-instances VPN-A protocols pim rp static address 10.255.182.144
set routing-instances VPN-A protocols pim interface lo0.1 mode sparse-dense
set routing-instances VPN-A protocols pim interface fe-1/1/2.0 mode sparse-dense
set routing-instances VPN-A protocols mvpn family inet autodiscovery-only intra-as inclusive
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.182.133
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family inet-mdt signaling
set protocols bgp group ibgp neighbor 10.255.182.142
set protocols ospf traffic-engineering
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 pim rp local address 10.255.182.133
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept

```

Router PE2

```

set interfaces so-0/0/1 unit 0 family inet address 10.10.20.1/30
set interfaces so-0/0/3 unit 0 family inet address 10.111.10.2/30
set interfaces so-0/0/3 unit 0 family iso
set interfaces so-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.182.142/32 primary

```

```
set interfaces lo0 unit 1 family inet address 10.10.47.101/32
set routing-options router-id 10.255.182.142
set routing-options route-distinguisher-id 10.255.182.142
set routing-options autonomous-system 100
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface so-0/0/1.0
set routing-instances VPN-A interface lo0.1
set routing-instances VPN-A provider-tunnel pim-ssm family inet group-address 232.1.1.1
set routing-instances VPN-A provider-tunnel mdt threshold group 224.1.1.0/24 source
10.240.0.242/32 rate 10
set routing-instances VPN-A provider-tunnel mdt tunnel-limit 20
set routing-instances VPN-A provider-tunnel mdt group-range 232.1.1.3/32
set routing-instances VPN-A vrf-target target:100:10
set routing-instances VPN-A vrf-table-label
set routing-instances VPN-A routing-options graceful-restart
set routing-instances VPN-A protocols ospf area 0.0.0.0 interface all
set routing-instances VPN-A protocols ospf export bgp-to-ospf
set routing-instances VPN-A protocols pim mvpn family inet autodiscovery inet-mdt
set routing-instances VPN-A protocols pim mvpn family inet6 disable
set routing-instances VPN-A protocols pim rp static address 10.255.182.144
set routing-instances VPN-A protocols pim interface lo0.1 mode sparse-dense
set routing-instances VPN-A protocols pim interface so-0/0/1.0 mode sparse-dense
set routing-instances VPN-A protocols mvpn family inet autodiscovery-only intra-as inclusive
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.182.142
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family inet-mdt signaling
set protocols bgp group ibgp neighbor 10.255.182.133
set protocols ospf traffic-engineering
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 pim rp static address 10.255.182.133
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept
```

Router 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 the [Junos OS CLI User Guide](#).

To configure Router PE1:

1. Configure the router interfaces, enabling IPv4 traffic.

Also enable MPLS on the interface facing Router PE2.

The lo0.1 interface is for the VPN-A routing instance.

```
[edit interfaces]
user@PE1# set so-0/0/3 unit 0 family inet address 10.111.10.1/30
user@PE1# set so-0/0/3 unit 0 family mpls
user@PE1# set fe-1/1/2 unit 0 family inet address 10.10.10.1/30
user@PE1# set lo0 unit 0 family inet address 10.255.182.133/32 primary
user@PE1# set lo0 unit 1 family inet address 10.10.47.100/32
```

2. Configure a routing policy to export BGP routes from the routing table into OSPF.

```
[edit policy-options policy-statement bgp-to-ospf]
user@PE1# set from protocol bgp
user@PE1# set then accept
```

3. Configure the router ID, route distinguisher, and autonomous system number.

```
[edit routing-options]
user@PE1# set router-id 10.255.182.133
user@PE1# set route-distinguisher-id 10.255.182.133
user@PE1# set autonomous-system 100
```

4. Configure the protocols that need to run in the main routing instance to enable MPLS, BGP, the IGP, VPNs, and PIM sparse mode.

```
[edit protocols ]
user@PE1# set mpls interface all
```

```

user@PE1# set mpls interface fxp0.0 disable
user@PE1# set bgp group ibgp type internal
user@PE1# set bgp group ibgp local-address 10.255.182.133
user@PE1# set bgp group ibgp family inet-vpn unicast
user@PE1# set bgp group ibgp neighbor 10.255.182.142
user@PE1# set ospf traffic-engineering
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 ldp interface all
user@PE1# set pim rp local address 10.255.182.133
user@PE1# set pim interface all mode sparse
user@PE1# set pim interface all version 2
user@PE1# set pim interface fxp0.0 disable

```

5. Create the customer VRF routing instance.

```

[edit routing-instances VPN-A]
user@PE1# set instance-type vrf
user@PE1# set interface fe-1/1/2.0
user@PE1# set interface lo0.1
user@PE1# set vrf-target target:100:10
user@PE1# set vrf-table-label
user@PE1# set protocols ospf area 0.0.0.0 interface all
user@PE1# set protocols ospf export bgp-to-ospf
user@PE1# set protocols pim rp static address 10.255.182.144
user@PE1# set protocols pim interface lo0.1 mode sparse-dense
user@PE1# set protocols pim interface fe-1/1/2.0 mode sparse-dense

```

6. Configure the draft-rosen 7 control plane, and specify IPv4 traffic to be carried in the provider tunnel.

```

[edit routing-instances VPN-A]
user@PE1# set provider-tunnel pim-ssm family inet group-address 232.1.1.1
user@PE1# set protocols pim mvpn family inet autodiscovery inet-mdt
user@PE1# set protocols pim mvpn family inet6 disable
user@PE1# set protocols mvpn family inet autodiscovery-only intra-as inclusive
[edit protocols bgp group ibgp]
user@PE1# set family inet-mdt signaling

```

7. (Optional) Configure a data MDT tunnel.

```
[edit routing-instances VPN-A]
user@PE1# set provider-tunnel mdt threshold group 224.1.1.0/24 source 10.240.0.242/32 rate 10
user@PE1# set provider-tunnel mdt tunnel-limit 20
user@PE1# set provider-tunnel mdt group-range 232.1.1.3/32
```

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
lo0 {
  unit 0 {
    family inet {
      address 10.255.182.133/32 {
        primary;
      }
    }
  }
  unit 1 {
    family inet {
      address 10.10.47.100/32;
    }
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 10.111.10.1/30;
    }
    family mpls;
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.10.10.1/30;
    }
  }
}
```

```
}  
}
```

```
user@PE1# show policy-options  
policy-statement bgp-to-ospf {  
    from protocol bgp;  
    then accept;  
}
```

```
user@PE1# show protocols  
mpls {  
    ipv6-tunneling;  
    interface all;  
    interface fxp0.0 {  
        disable;  
    }  
}  
bgp {  
    group ibgp {  
        type internal;  
        local-address 10.255.182.133;  
        family inet-vpn {  
            unicast;  
        }  
        family inet-mdt {  
            signaling;  
        }  
        neighbor 10.255.182.142;  
    }  
}  
ospf {  
    traffic-engineering;  
    area 0.0.0.0 {  
        interface all;  
        interface fxp0.0 {  
            disable;  
        }  
    }  
}  
ldp {
```



```

    interface all;
}
pim {
    rp {
        local {
            address 10.255.182.133;
        }
    }
    interface all {
        mode sparse;
        version 2;
    }
    interface fxp0.0 {
        disable;
    }
}
}

```

```

user@PE1# show routing-instances
VPN-A {
    instance-type vrf;
    interface fe-1/1/2.0;
    interface lo0.1;
    provider-tunnel {
        pim-ssm {
            family {
                inet {
                    group-address 232.1.1.1;
                }
            }
        }
        mdt {
            threshold {
                group 224.1.1.0/24 {
                    source 10.240.0.242/32 {
                        rate 10;
                    }
                }
            }
        }
        tunnel-limit 20;
        group-range 232.1.1.3/32;
    }
}

```

```
}
vrf-target target:100:10;
vrf-table-label;
protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.0 {
      interface all;
    }
  }
  pim {
    mvpn {
      family {
        inet {
          autodiscovery {
            inet-mdt;
          }
        }
        inet6 {
          disable;
        }
      }
    }
    rp {
      static {
        address 10.255.182.144;
      }
    }
    interface lo0.1 {
      mode sparse-dense;
    }
    interface fe-1/1/2.0 {
      mode sparse-dense;
    }
  }
  mvpn {
    family {
      inet {
        autodiscovery-only {
          intra-as {
            inclusive;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}
}
}

```

```

user@PE1# show routing-options
route-distinguisher-id 10.255.182.133;
autonomous-system 100;
router-id 10.255.182.133;

```

If you are done configuring the router, enter `commit` from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and IP addresses.

CE Device Configuration

IN THIS SECTION

- [CLI Quick Configuration | 641](#)
- [Device CE1 | 642](#)
- [Results | 643](#)

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 fe-0/1/0 unit 0 family inet address 10.10.10.2/30
set interfaces lo0 unit 0 family inet address 10.255.182.144/32 primary
set routing-options router-id 10.255.182.144
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 local address 10.255.182.144

```

```
set protocols pim interface all mode sparse-dense
set protocols pim interface fxp0.0 disable
```

Device CE2

```
set interfaces so-0/0/1 unit 0 family inet address 10.10.20.2/30
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.182.140/32 primary
set routing-options router-id 10.255.182.140
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 10.255.182.144
set protocols pim interface all mode sparse-dense
set protocols pim interface fxp0.0 disable
```

Device CE1

Step-by-Step Procedure

To configure Device CE1:

1. Configure the router interfaces, enabling IPv4 and IPv6 traffic.

```
[edit interfaces]
user@CE1# set fe-0/1/0 unit 0 family inet address 10.10.10.2/30
user@CE1# set lo0 unit 0 family inet address 10.255.182.144/32 primary
```

2. Configure the router ID.

```
[edit routing-options]
user@CE1# set router-id 10.255.182.144
```

3. Configure the protocols that need to run on the CE device to enable OSPF (for IPv4) and PIM sparse-dense mode.

```
[edit protocols]
user@CE1# set ospf area 0.0.0.0 interface all
user@CE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@CE1# set pim rp local address 10.255.182.144
```

```
user@CE1# set pim interface all mode sparse-dense
user@CE1# set pim interface fxp0.0 disable
```

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 configuration instructions in this example to correct it.

```
user@CE1# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.10.10.2/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.182.144/32 {
        primary;
      }
    }
  }
}
```

```
user@CE1# show protocols
ospf {
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
pim {
  rp {
```

```
    local {
        address 10.255.182.144;
    }
}
interface all {
    mode sparse-dense;
}
interface fxp0.0 {
    disable;
}
}
```

```
user@CE1# show routing-options
router-id 10.255.182.144;
```

If you are done configuring the router, enter `commit` from configuration mode.

Repeat the procedure for Device CE2, using the appropriate interface names and IP addresses.

Verification

IN THIS SECTION

- [Verifying Tunnel Encapsulation | 644](#)
- [Verifying PIM Neighbors | 645](#)
- [Verifying the Provider Tunnel and Control Plane | 646](#)
- [Checking Routes | 646](#)
- [Verifying MDT Tunnels | 647](#)

Confirm that the configuration is working properly.

Verifying Tunnel Encapsulation

Purpose

Verify that PIM multicast tunnel (mt) encapsulation and deencapsulation interfaces come up.

Action

```
user@PE1> show pim interfaces instance VPN-A
```

```
Instance: PIM.VPN-A
```

Name	Stat	Mode	IP V	State	NbrCnt	JoinCnt(sg)	JoinCnt(*g)	DR address
fe-1/1/2.0	Up	SparseDense	4 2	NotDR	1	1	1	10.10.10.2
lo0.1	Up	SparseDense	4 2	DR	0	0	0	10.10.47.100
lsi.2304	Up	SparseDense	4 2	P2P	0	0	0	
mt-0/3/0.32769	Up	SparseDense	4 2	P2P	0	0	0	
mt-1/2/0.1081344	Up	SparseDense	4 2	P2P	0	0	0	
mt-1/2/0.32768	Up	SparseDense	4 2	P2P	1	0	0	
pe-0/3/0.32770	Up	Sparse	4 2	P2P	0	0	0	

Meaning

The multicast tunnel interface that is used for encapsulation, `mt-[xxxx]`, is in the range from 32,768 through 49,151. The interface `mt-[yyyy]`, used for de-encapsulation, is in the range from 1,081,344 through 1,107,827. PIM runs only on the encapsulation interface. The de-encapsulation interface populates downstream interface information.

Verifying PIM Neighbors

Purpose

Verify that PIM neighborhood is established over the multicast tunnel interface.

Action

```
user@PE1> show pim neighbors instance VPN-A
```

```
Instance: PIM.VPN-A
```

```
B = Bidirectional Capable, G = Generation Identifier,
```

```
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
```

```
P = Hello Option DR Priority, T = Tracking Bit
```

Interface	IP V	Mode	Option	Uptime	Neighbor	addr
fe-1/1/2.0	4	2	HPLGT	00:29:35	10.10.10.2	
mt-1/2/0.32768	4	2	HPLGT	00:28:32	10.10.47.101	

Meaning

When the neighbor address is listed and the uptime is incrementing, it means that PIM neighborship is established over the multicast tunnel interface.

Verifying the Provider Tunnel and Control Plane

Purpose

Confirm that the provider tunnel and control-plane protocols are correct.

Action

```
user@PE1> show pim mvpn
Instance      Family      VPN-Group    Mode      Tunnel
PIM.VPN-A    INET        225.1.1.1    PIM-MVPN  PIM-SSM
```

Meaning

For draft-rosen, the MVPN mode appears in the output as PIM-MVPN.

Checking Routes

Purpose

Verify that traffic flows as expected.

Action

```
user@R1> show multicast route extensive instance VPN-A
Family: INET

Group: 224.1.1.1
  Source: 10.240.0.242/32
  Upstream interface: fe-1/1/2.0
  Downstream interface list:
    mt-1/2/0.32768
  Session description: NOB Cross media facilities
  Statistics: 92 kBps, 1001 pps, 1869820 packets
```



```

Next-hop ID: 1048581
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: 360 seconds
Wrong incoming interface notifications: 0

```

Meaning

For draft-rosen, the upstream protocol appears in the output as PIM.

Verifying MDT Tunnels

Purpose

Verify that both default and data MDT tunnels are correct.

Action

```

user@PE1> show pim mdt instance VPN-A
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.182.133
Default tunnel interface: mt-1/2/0.32769
Default tunnel source: 0.0.0.0

C-group address   C-source address   P-group address   Data tunnel interface
224.1.1.1         10.240.0.242       232.1.1.3         mt-0/3/0.32771

Instance: PIM.VPN-A
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.182.142
Default tunnel interface: mt-1/2/0.1081345
Default tunnel source: 0.0.0.0

```

RELATED DOCUMENTATION

[Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs | 649](#)

[Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode | 669](#)

Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs

IN THIS SECTION

- [Understanding Source-Specific Multicast VPNs | 648](#)
- [Draft-Rosen 7 Multicast VPN Control Plane | 649](#)
- [Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs | 649](#)

Understanding Source-Specific Multicast VPNs

A draft-rosen MVPN with service provider tunnels operating in SSM mode uses BGP signaling for autodiscovery of the PE routers. These MVPNs are also referred to as Draft Rosen 7.

Each PE sends an MDT subsequent address family identifier (MDT-SAFI) BGP network layer reachability information (NLRI) advertisement. The advertisement contains the following information:

- Route distinguisher
- Unicast address of the PE router to which the source site is attached (usually the loopback)
- Multicast group address
- Route target extended community attribute

Each remote PE router imports the MDT-SAFI advertisements from each of the other PE routers if the route target matches. Each PE router then joins the (S,G) tree rooted at each of the other PE routers.

After a PE router discovers the other PE routers, the source and group are bound to the VPN routing and forwarding (VRF) through the multicast tunnel de-encapsulation interface.

A draft-rosen MVPN with service provider tunnels operating in any-source multicast sparse-mode uses a shared tree and rendezvous point (RP) for autodiscovery of the PE routers. The PE that is the source of the multicast group encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router. The RP then builds a shortest-path tree (SPT) toward the source PE.

The remote PE that acts as a receiver for the MDT multicast group sends (*,G) join messages toward the RP and joins the distribution tree for that group.

Draft-Rosen 7 Multicast VPN Control Plane

The control plane of a draft-rosen MVPN with service provider tunnels operating in SSM mode must be configured to support autodiscovery.

After the PE routers are discovered, PIM is notified of the multicast source and group addresses. PIM binds the (S,G) state to the multicast tunnel (**mt**) interface and sends a join message for that group.

Autodiscovery for a draft-rosen MVPN with service provider tunnels operating in SSM mode uses some of the facilities of the BGP-based MVPN control plane software module. Therefore, the BGP-based MVPN control plane must be enabled. The BGP-based MVPN control plane can be enabled for autodiscovery only.

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

IN THIS SECTION

- [Requirements | 649](#)
- [Overview | 650](#)
- [Configuration | 652](#)
- [Verification | 660](#)

This example shows how to configure a draft-rosen Layer 3 VPN operating in source-specific multicast (SSM) mode. This example is based on the Junos OS implementation of the IETF Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later
- Make sure that the routing devices support multicast tunnel (**mt**) interfaces.

A tunnel-capable PIC supports a maximum of 512 multicast tunnel interfaces. Both default and data MDTs contribute to this total. The default MDT uses two multicast tunnel interfaces (one for encapsulation and one for de-encapsulation). To enable an M Series or T Series router to support more than 512 multicast tunnel interfaces, another tunnel-capable PIC is required. See "[Tunnel](#)

[Services PICs and Multicast](#) on page 315 and ["Load Balancing Multicast Tunnel Interfaces Among Available PICs"](#) on page 611.



NOTE: In Junos OS Release 17.3R1, the `pim-ssm` hierarchy was moved from `provider-tunnel` to the `provider-tunnel family inet` and `provider-tunnel family inet6` 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.

Overview

IN THIS SECTION

- [Topology | 651](#)

The IETF Internet draft `draft-rosen-vpn-mcast-07.txt` introduced the ability to configure the provider network to operate in SSM mode. When a draft-rosen multicast VPN is used over an SSM provider core, there are no PIM RPs to provide rendezvous and autodiscovery between PE routers. Therefore, `draft-rosen-vpn-mcast-07` specifies the use of a BGP network layer reachability information (NLRI), called MDT subaddress family identifier information (MDT-SAFI) to facilitate autodiscovery of PEs by other PEs. MDT-SAFI updates are BGP messages distributed between intra-AS internal BGP peer PEs. Thus, receipt of an MDT-SAFI update enables a PE to autodiscover the identity of other PEs with sites for a given VPN and the default MDT (S,G) routes to join for each. Autodiscovery provides the next-hop address of each PE, and the VPN group address for the tunnel rooted at that PE for the given route distinguisher (RD) and route-target extended community attribute.

This example includes the following configuration options to enable draft-rosen SSM:

- **`protocols bgp group group-name family inet-mdt signaling`**—Enables MDT-SAFI signaling in BGP.
- **`routing-instance instance-name protocols mvpn family inet autodiscovery-only intra-as inclusive`**—Enables the multicast VPN to use the MDT-SAFI autodiscovery NLRI.
- **`routing-instance instance-name protocols pim mvpn`**—Specifies the SSM control plane. When `pim mvpn` is configured for a VRF, the VPN group address must be specified with the `provider-tunnel pim-ssm group-address` statement.
- **`routing-instance instance-name protocols pim mvpn family inet autodiscovery inet-mdt`**—Enables PIM to learn about neighbors from the MDT-SAFI autodiscovery NLRI.

- **routing-instance *instance-name* provider-tunnel family inet pim-ssm group-address *multicast-address***
—Configures the provider tunnel that serves as the control plane and enables the provider tunnel to have a static group address. Unlike draft-rosen multicast VPNs with ASM provider cores, the SSM configuration does not require that each PE for a VPN use the same group address. This is because the rendezvous point assignment and autodiscovery are not accomplished over the default MDT tunnels for the group. Thus, you can configure some or all PEs in a VPN to use a different group, but the same group cannot be used in different VPNs on the same PE router.
- **routing-instances *ce1* vrf-target target:100:1**—Configures the VRF export policy. When you configure draft-rosen multicast VPNs with provider tunnels 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.




NOTE: When you configure draft-rosen multicast VPNs with provider tunnels 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.

Topology

Figure 86 on page 652 shows the topology for this example.

Figure 86: SSM for Draft-Rosen Multicast VPNs Topology



Configuration

IN THIS SECTION

- Procedure | 653
- Interface Configuration | 655
- Multicast Group Management | 656

- [MPLS Signaling Protocol and MPLS LSPs | 656](#)
- [BGP | 657](#)
- [Interior Gateway Protocol | 658](#)
- [PIM | 659](#)
- [Routing Instance | 659](#)

Procedure

CLI Quick Configuration

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

```

set interfaces so-0/0/0 description "T0 P1_P1"
set interfaces so-0/0/0 unit 0 description "to P1 (provider router) so-0/0/0.0"
set interfaces so-0/0/0 unit 0 family inet address 10.0.1.1/30
set interfaces so-0/0/0 unit 0 family iso
set interfaces so-0/0/0 unit 0 family mpls
set interfaces so-0/0/1 description "T0 PE2"
set interfaces so-0/0/1 unit 0 description "to PE2 (PE router) so-0/0/1.0"
set interfaces so-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces so-0/0/1 unit 0 family iso
set interfaces so-0/0/1 unit 0 family mpls
set interfaces fe-0/1/1 description "T0 CE1"
set interfaces fe-0/1/1 unit 0 description "to CE router fe-0/1/1.0"
set interfaces fe-0/1/1 unit 0 family inet address 10.0.3.1/30
set interfaces lo0 unit 0 description "PE1 (this PE router) Loopback"
set interfaces lo0 unit 1 family inet address 10.1.1.0/32
set routing-options autonomous-system 65200
set protocols igmp query-interval 2
set protocols igmp query-response-interval 1
set protocols igmp query-last-member-interval 1
set protocols igmp interface all immediate-leave
set protocols igmp interface fxp0.0 disable
set protocols rsvp interface all
set protocols rsvp interface so-0/0/0.0
set protocols rsvp interface so-0/0/1.0

```

```
set protocols mpls label-switched-path PE1-to-PE2 to 10.255.14.217
set protocols mpls label-switched-path PE1-to-PE2 primary PE1_PE2_prime
set protocols mpls label-switched-path PE1-to-P1 to 10.255.14.218
set protocols mpls label-switched-path PE1-to-P1 primary PE1_P1_prime
set protocols mpls path PE1_P1_prime 10.0.1.2
set protocols mpls path PE1_PE2_prime 10.0.2.2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.14.216
set protocols bgp group int family inet unicast
set protocols bgp group int family inet-vpn unicast
set protocols bgp group int family inet-vpn multicast
set protocols bgp group int family inet-mdt signaling
set protocols bgp group int neighbor 10.255.14.218
set protocols bgp group int neighbor 10.255.14.217
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 so-0/0/0.0 metric 10
set protocols ospf area 0.0.0.0 interface so-0/0/1.0 metric 10
set protocols pim assert-timeout 5
set protocols pim join-prune-timeout 210
set protocols pim rp bootstrap-priority 10
set protocols pim rp local address 10.255.14.216
set protocols pim interface lo0.0
set protocols pim interface all hello-interval 1
set protocols pim interface fxp0.0 disable
set policy-options policy-statement bgp_ospf term 1 from protocol bgp
set policy-options policy-statement bgp_ospf term 1 then accept
set routing-instances ce1 instance-type vrf
set routing-instances ce1 interface fe-0/1/1.0
set routing-instances ce1 interface lo0.1
set routing-instances ce1 route-distinguisher 10:0
set routing-instances ce1 provider-tunnel pim-ssm group-address 232.1.1.1
set routing-instances ce1 vrf-target target:100:1
set routing-instances ce1 protocols ospf export bgp_ospf
set routing-instances ce1 protocols ospf sham-link local 01.1.1.0
set routing-instances ce1 protocols ospf area 0.0.0.0 sham-link-remote 10.1.1.1
set routing-instances ce1 protocols ospf area 0.0.0.0 sham-link-remote 10.1.1.2
set routing-instances ce1 protocols ospf area 0.0.0.0 interface lo0.1
set routing-instances ce1 protocols ospf area 0.0.0.0 interface fe-0/1/1.0 metric 10
set routing-instances ce1 protocols pim mvpn family inet autodiscovery inet-mdt
set routing-instances ce1 protocols pim interface lo0.1
```



```

set routing-instances ce1 protocols pim interface fe-0/1/1.0 priority 100
set routing-instances ce1 protocols pim interface fe-0/1/1.0 hello-interval 1
set routing-instances ce1 protocols mvpn family inet autodiscovery-only intra-as inclusive

```

Interface Configuration

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 [Junos OS CLI User Guide](#).

To configure the interfaces on one PE router:

1. Configure PE1's interface to the provider router.

```

[edit interfaces so-0/0/0]
user@host# set description "T0 P1"
user@host# set unit 0 description "to P1 (provider router, 10.255.14.218 ) so-0/0/0.0"
user@host# set unit 0 family inet address 10.0.1.1/30
user@host# set unit 0 family iso
user@host# set unit 0 family mpls

```

2. Configure PE1's interface to PE2.

```

[edit interfaces so-0/0/1]
user@host# set description "T0 PE2"
user@host# set unit 0 description "to PE2 (10.255.14.217) so-0/0/1.0"
user@host# set unit 0 family inet address 10.0.2.1/30
user@host# set unit 0 family iso
user@host# set unit 0 family mpls

```

3. Configure PE1's interface to CE1.

```

[edit interfaces fe-0/1/1]
user@host# set description "T0 CE1"
user@host# set unit 0 description "to CE1 (10.255.14.223) fe-0/1/1.0"
user@host# set unit 0 family inet address 10.0.3.1/30

```

```
user@host# set unit 0 family iso
user@host# set unit 0 family mpls
```

4. Configure PE1's loopback interface.

```
[edit interfaces lo0]
user@host# set unit 0 description "PE1 (this PE router, 10.255.14.216) Loopback"
user@host# set unit 1 family inet address 10.1.1.0/32
```

Multicast Group Management

Step-by-Step Procedure

To configure multicast group management:

1. Configure the IGMP interfaces.

```
[edit protocols igmp]
user@host# set interface all immediate-leave
user@host# set interface fxp0.0 disable
```

2. Configure the IGMP settings.

```
[edit protocols igmp]
user@host# set query-interval 2
user@host# set query-response-interval 1
user@host# set query-last-member-interval 1
```

MPLS Signaling Protocol and MPLS LSPs

Step-by-Step Procedure

To configure the MPLS signaling protocol and MPLS LSPs:

1. Configure RSVP signaling among this PE router (PE1), the other PE router (PE2), and the provider router (P1).

```
[edit protocols rsvp]
user@host# set interface so-0/0/0.0
user@host# set interface so-0/0/1.0
```

2. Configure MPLS LSPs.

```
[edit protocols mpls]
user@host# set label-switched-path pe1-to-pe2 to 10.255.14.217
user@host# set label-switched-path pe1-to-pe2 primary pe1_pe2_prime
user@host# set label-switched-path pe1-to-p1 to 10.255.14.218
user@host# set label-switched-path pe1-to-p1 primary pe1_p1_prime
user@host# set path pe1_p1_prime 10.0.1.2
user@host# set path pe1_pe2_prime 10.0.2.2
user@host# set interface all
user@host# set interface fxp0.0 disable
```

BGP

Step-by-Step Procedure

To configure BGP:

1. Configure the AS number. In this example, both of the PE routers and the provider router are in AS 65200.

```
[edit]
user@host# set routing-options autonomous-system 65200
```

2. Configure the internal BGP full mesh with the PE2 and P1 routers.

```
[edit protocols bgp group int]
user@host# set type internal
user@host# set local-address 10.255.14.216
user@host# set family inet unicast
```

```
user@host# set neighbor 10.255.14.218
user@host# set neighbor 10.255.14.217
```

3. Enable MDT-SAFI NLRI control plane messages.

```
[edit protocols bgp group int]
user@host# set family inet-mdt signaling
```

4. Enable BGP to carry Layer 3 VPN NLRI for the IPv4 address family.

```
[edit protocols bgp group int]
user@host# set family inet-vpn unicast
user@host# set family inet-vpn multicast
```

5. Configure BGP export policy.

```
[edit policy-options]
user@host# set policy-statement bgp_ospf term 1 from protocol bgp
user@host# set policy-statement bgp_ospf term 1 then accept
```

Interior Gateway Protocol

Step-by-Step Procedure

To configure the interior gateway protocol:

1. Configure the OSPF interfaces.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface lo0.0 passive
user@host# set area 0.0.0.0 interface so-0/0/0.0 metric 10
user@host# set area 0.0.0.0 interface so-0/0/1.0 metric 10
```

2. Enable traffic engineering.

```
[edit protocols ospf]
user@host# set traffic-engineering
```

PIM

Step-by-Step Procedure

To configure PIM:

1. Configure timeout periods and the RP. Local RP configuration makes PE1 a statically defined RP.

```
[edit protocols pim]
user@host# set assert-timeout 5
user@host# set join-prune-timeout 210
user@host# set rp bootstrap-priority 10
user@host# set rp local address 10.255.14.216
```

2. Configure the PIM interfaces.

```
[edit protocols pim]
user@host# set interface lo0.0
user@host# set interface all hello-interval 1
user@host# set interface fxp0.0 disable
```

Routing Instance

Step-by-Step Procedure

To configure the routing instance between PE1 and CE1:

1. Configure the basic routing instance.

```
[edit routing-instances ce1]
user@host# set instance-type vrf
user@host# set interface fe-0/1/1.0
user@host# set interface lo0.1
user@host# set route-distinguisher 10:0
user@host# set vrf-target target:100:1
```

2. Configure the SSM provider tunnel.

```
[edit routing-instances ce1]
user@host# set provider-tunnel family inet pim-ssm group-address (Routing Instances) 232.1.1.1
```

3. Configure OSPF in the routing instance.

```
[edit routing-instances ce1 protocols ospf]
user@host# set export bgp_ospf
user@host# set sham-link local 10.1.1.0
user@host# set area 0.0.0.0 sham-link-remote 10.1.1.1
user@host# set area 0.0.0.0 sham-link-remote 10.1.1.2
user@host# set area 0.0.0.0 interface lo0.1
user@host# set area 0.0.0.0 interface fe-0/1/1.0 metric 10
```

4. Configure PIM in the routing instance.

```
[edit routing-instances ce1 protocols pim]
user@host# set interface lo0.1
user@host# set interface fe-0/1/1.0 priority 100
user@host# set interface fe-0/1/1.0 hello-interval 1
```

5. Configure draft-rosen VPN autodiscovery for provider tunnels operating in SSM mode.

```
[edit routing-instances ce1 protocols pim ]
user@host# set mvpn family inet autodiscovery inet-mdt
```

6. Configure the BGP-based MVPN control plane to provide signaling only for autodiscovery and not for PIM operations.

```
[edit routing-instances ce1 protocols mvpn family inet]
user@host# set autodiscovery-only intra-as inclusive
```

Verification

You can monitor the operation of the routing instance by running the `show route table ce1.mdt.0` command.

You can manage the group-instance mapping for local SSM tunnel roots by running the `show pim mvpn` command.

The `show pim mdt` command shows the tunnel type and source PE address for each outgoing and incoming MDT. In addition, because each PE might have its own default MDT group address, one incoming entry is shown for each remote PE. Outgoing data MDTs are shown after the outgoing default MDT. Incoming data MDTs are shown after all incoming default MDTs.

For troubleshooting, you can configure tracing operations for all of the protocols.

SEE ALSO

[Draft-Rosen Multicast VPNs Overview | 610](#)

[Understanding Data MDTs | 661](#)

[Data MDT Characteristics | 684](#)

RELATED DOCUMENTATION

[Example: Configuring Any-Source Draft-Rosen 6 Multicast VPNs | 611](#)

Understanding Data MDTs

In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, the VPN is multicast-enabled and configured to use the Protocol Independent Multicast (PIM) protocol within the VPN and within the service provider (SP) network. A multicast-enabled VPN routing and forwarding (VRF) instance corresponds to a multicast domain (MD), and a PE router attached to a particular VRF instance is said to belong to the corresponding MD. For each MD there is a default multicast distribution tree (MDT) through the SP backbone, which connects all of the PE routers belonging to that MD. Any PE router configured with a default MDT group address can be the multicast source of one default MDT.

To provide optimal multicast routing, you can configure the PE routers so that when the multicast source within a site exceeds a traffic rate threshold, the PE router to which the source site is attached creates a new data MDT and advertises the new MDT group address. An advertisement of a new MDT group address is sent in a User Datagram Protocol (UDP) type-length-value (TLV) packet called an *MDT join TLV*. The MDT join TLV identifies the source and group pair (S,G) in the VRF instance as well as the new data MDT group address used in the provider space. The PE router to which the source site is attached sends the MDT join TLV over the default MDT for that VRF instance every 60 seconds as long as the source is active.

All PE routers in the VRF instance receive the MDT join TLV because it is sent over the default MDT, but not all the PE routers join the new data MDT group:

- PE routers connected to receivers in the VRF instance for the current multicast group cache the contents of the MDT join TLV, adding a 180-second timeout value to the cache entry, and also join the new data MDT group.
- PE routers not connected to receivers listed in the VRF instance for the current multicast group also cache the contents of the MDT join TLV, adding a 180-second timeout value to the cache entry, but do not join the new data MDT group at this time.

After the source PE stops sending the multicast traffic stream over the default MDT and uses the new MDT instead, only the PE routers that join the new group receive the multicast traffic for that group.

When a remote PE router joins the new data MDT group, it sends a PIM join message for the new group directly to the source PE router from the remote PE routers by means of a PIM (S,G) join.

If a PE router that has not yet joined the new data MDT group receives a PIM join message for a new receiver for which (S,G) traffic is already flowing over the data MDT in the provider core, then that PE router can obtain the new group address from its cache and can join the data MDT immediately without waiting up to 59 seconds for the next data MDT advertisement.

When the PE router to which the source site is attached sends a subsequent MDT join TLV for the VRF instance over the default MDT, any existing cache entries for that VRF instance are simply refreshed with a timeout value of 180 seconds.

To display the information cached from MDT join TLV packets received by all PE routers in a PIM-enabled VRF instance, use the **show pim mdt data-mdt-joins** *operational mode command*.

The source PE router starts encapsulating the multicast traffic for the VRF instance using the new data MDT group after 3 seconds, allowing time for the remote PE routers to join the new group. The source PE router then halts the flow of multicast packets over the default MDT, and the packet flow for the VRF instance source shifts to the newly created data MDT.

The PE router monitors the traffic rate during its periodic statistics-collection cycles. If the traffic rate drops below the threshold or the source stops sending multicast traffic, the PE router to which the source site is attached stops announcing the MDT join TLVs and switches back to sending on the default MDT for that VRF instance.

RELATED DOCUMENTATION

show pim mdt data-mdt-joins

[CLI Explorer](#)

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

IN THIS SECTION

- Requirements | 663
- Overview | 663
- Configuration | 666
- Verification | 668

This example shows how to configure data multicast distribution trees (MDTs) in a draft-rosen Layer 3 VPN operating in any-source multicast (ASM) mode. This example is based on the Junos OS implementation of RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and on section 2 of the IETF Internet draft draft-rosen-vpn-mcast-06.txt, *Multicast in MPLS/BGP VPNs* (expired April 2004).

Requirements

Before you begin:

- Configure the draft-rosen multicast over Layer 3 VPN scenario.
- Make sure that the routing devices support multicast tunnel (**mt**) interfaces.

A tunnel-capable PIC supports a maximum of 512 multicast tunnel interfaces. Both default and data MDTs contribute to this total. The default MDT uses two multicast tunnel interfaces (one for encapsulation and one for de-encapsulation). To enable an M Series or T Series router to support more than 512 multicast tunnel interfaces, another tunnel-capable PIC is required. See [""Tunnel Services PICs and Multicast" on page 315"](#) and [""Load Balancing Multicast Tunnel Interfaces Among Available PICs" on page 625"](#) in the Multicast Protocols User Guide..

Overview

IN THIS SECTION

- Topology | 665

By using data multicast distribution trees (MDTs) in a Layer 3 VPN, you can prevent multicast packets from being flooded unnecessarily to specified provider edge (PE) routers within a VPN group. This option is primarily useful for PE routers in your Layer 3 VPN multicast network that have no receivers for the multicast traffic from a particular source.

When a PE router that is directly connected to the multicast source (also called the *source PE*) receives Layer 3 VPN multicast traffic that exceeds a configured threshold, a new data MDT tunnel is established between the PE router connected to the source site and its remote PE router neighbors.

The source PE advertises the new data MDT group as long as the source is active. The periodic announcement is sent over the default MDT for the VRF. Because the data MDT announcement is sent over the default tunnel, all the PE routers receive the announcement.

Neighbors that do not have receivers for the multicast traffic cache the advertisement of the new data MDT group but ignore the new tunnel. Neighbors that do have receivers for the multicast traffic cache the advertisement of the new data MDT group and also send a PIM join message for the new group.

The source PE encapsulates the VRF multicast traffic using the new data MDT group and stops the packet flow over the default multicast tree. If the multicast traffic level drops back below the threshold, the data MDT is torn down automatically and traffic flows back across the default multicast tree.

If a PE router that has not yet joined the new data MDT group receives a PIM join message for a new receiver for which (S,G) traffic is already flowing over the data MDT in the provider core, then that PE router can obtain the new group address from its cache and can join the data-MDT immediately without waiting up to 59 seconds for the next data MDT advertisement.

By default, automatic creation of data MDTs is disabled.

For a rosen 6 MVPN—a draft-rosen multicast VPN with provider tunnels operating in ASM mode—you configure data MDT creation for a tunnel multicast group by including statements under the PIM protocol configuration for the VRF instance associated with the multicast group. Because data MDTs apply to VPNs and VRF routing instances, you cannot configure MDT statements in the master routing instance.

This example includes the following configuration options:

- **group**—Specifies the multicast group address to which the threshold applies. This could be a well-known address for a certain type of multicast traffic.

The group address can be explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). Explicit and prefix address forms can be combined if they do not overlap. Overlapping configurations, in which prefix and more explicit address forms are used for the same source or group address, are not supported.

- **group-range**—Specifies the multicast group IP address range used when a new data MDT needs to be initiated on the PE router. For each new data MDT, one address is automatically selected from the configured group range.

The PE router implementing data MDTs for a local multicast source must be configured with a range of multicast group addresses. Group addresses that fall within the configured range are used in the join messages for the data MDTs created in this VRF instance. Any multicast address range can be used as the multicast prefix. However, the group address range cannot overlap the default MDT group address configured for any VPN on the router. If you configure overlapping group addresses, the configuration commit operation fails.

- **pim**—Supports data MDTs for service provider tunnels operating in any-source multicast mode.
- **rate**—Specifies the data rate that initiates the creation of data MDTs. When the source traffic in the VRF exceeds the configured data rate, a new tunnel is created. The range is from 10 kilobits per second (Kbps), the default, to 1 gigabit per second (Gbps, equivalent to 1,000,000 Kbps).
- **source**—Specifies the unicast address of the source of the multicast traffic. It can be a source locally attached to or reached through the PE router. A group can have more than one source.

The source address can be explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). Explicit and prefix address forms can be combined if they do not overlap. Overlapping configurations, in which prefix and more explicit address forms are used for the same source or group address, are not supported.

- **threshold**—Associates a rate with a group and a source. The PE router implementing data MDTs for a local multicast source must establish a data MDT-creation threshold for a multicast group and source.

When the traffic stops or the rate falls below the threshold value, the source PE router switches back to the default MDT.

- **tunnel-limit**—Specifies the maximum number of data MDTs that can be created for a single routing instance. The PE router implementing a data MDT for a local multicast source must establish a limit for the number of data MDTs created in this VRF instance. If the limit is 0 (the default), then no data MDTs are created for this VRF instance.

If the number of data MDT tunnels exceeds the maximum configured tunnel limit for the VRF, then no new tunnels are created. Traffic that exceeds the configured threshold is sent on the default MDT.

The valid range is from 0 through 1024 for a VRF instance. There is a limit of 8000 tunnels for all data MDTs in all VRF instances on a PE router.

Topology

Figure 87 on page 666 shows a default MDT.

Figure 87: Default MDT





Figure 88 on page 666 shows a data MDT.

Figure 88: Data MDT



Configuration

IN THIS SECTION

- Procedure | 667

Procedure

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.

```
[edit]
set routing-instances vpn-A protocols pim mdt group-range 227.0.0.0/8
set routing-instances vpn-A protocols pim mdt threshold group 224.4.4.4/32 source 10.10.20.43/32
rate 10
set routing-instances vpn-A protocols pim mdt tunnel-limit 10
```

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 [Junos OS CLI User Guide](#).

To configure a PE router attached to the VRF instance **vpn-A** in a PIM-ASM multicast VPN to initiate new data MDTs and provider tunnels for that VRF:

1. Configure the group range.

```
[edit]
user@host# edit routing-instances vpn-A protocols pim mdt
[edit routing-instances vpn-A protocols pim mdt]
user@host# set group-range 227.0.0.0/8
```

2. Configure a data MDT-creation threshold for a multicast group and source.

```
[edit routing-instances vpn-A protocols pim mdt]
user@host# set threshold group 224.4.4.4 source 10.10.20.43 rate 10
```

3. Configure a tunnel limit.

```
[edit routing-instances vpn-A protocols pim mdt]
user@host# set tunnel-limit 10
```

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

```
[edit routing-instances vpn-A protocols pim mdt]
user@host# commit
```

Verification

To display information about the default MDT and any data MDTs for the VRF instance **vpn-A**, use the **show pim mdt instance ce1 detail** operational mode command. This command displays either the outgoing tunnels (the tunnels initiated by the local PE router), the incoming tunnels (tunnels initiated by the remote PE routers), or both.

To display the data MDT group addresses cached by PE routers that participate in the VRF instance **vpn-A**, use the **show pim mdt data-mdt-joins instance vpn-A** operational mode command. The command displays the information cached from MDT join TLV packets received by all PE routers participating in the specified VRF instance.

You can trace the operation of data MDTs by including the **mdt detail** flag in the `[edit protocols pim traceoptions]` configuration. When this flag is set, all the **mt** interface-related activity is logged in trace files.

RELATED DOCUMENTATION

Introduction to Configuring Layer 3 VPNs

[Junos OS VPNs Library for Routing Devices](#)

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

IN THIS SECTION

- Requirements | 669
- Overview | 670
- Configuration | 676
- Verification | 681

This example shows how to configure data multicast distribution trees (MDTs) for a provider edge (PE) router attached to a VPN routing and forwarding (VRF) instance in a draft-rosen Layer 3 multicast VPN operating in source-specific multicast (SSM) mode. The example is based on the Junos OS implementation of RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and on section 7 of the IETF Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP IP VPNs*.

Requirements

Before you begin:

- Make sure that the routing devices support multicast tunnel (**mt**) interfaces.

A tunnel-capable PIC supports a maximum of 512 multicast tunnel interfaces. Both default and data MDTs contribute to this total. The default MDT uses two multicast tunnel interfaces (one for encapsulation and one for de-encapsulation). To enable an M Series or T Series router to support more than 512 multicast tunnel interfaces, another tunnel-capable PIC is required. See ["Tunnel Services PICs and Multicast" on page 315](#) and ["Load Balancing Multicast Tunnel Interfaces Among Available PICs" on page 625](#) in the Multicast Protocols User Guide.

- Make sure that the PE router has been configured for a draft-rosen Layer 3 multicast VPN operating in SSM mode in the provider core.

In this type of multicast VPN, PE routers discover one another by sending MDT subsequent address family identifier (MDT-SAFI) BGP network layer reachability information (NLRI) advertisements. Key configuration statements for the master instance are highlighted in [Table 18 on page 671](#). Key configuration statements for the VRF instance to which your PE router is attached are highlighted in [Table 19 on page 672](#). For complete configuration details, see ["Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs" on page 649](#) in the Multicast Protocols User Guide.

Overview

IN THIS SECTION

- [Topology | 676](#)

By using data MDTs in a Layer 3 VPN, you can prevent multicast packets from being flooded unnecessarily to specified provider edge (PE) routers within a VPN group. This option is primarily useful for PE routers in your Layer 3 VPN multicast network that have no receivers for the multicast traffic from a particular source.

- When a PE router that is directly connected to the multicast source (also called the *source PE*) receives Layer 3 VPN multicast traffic that exceeds a configured threshold, a new data MDT tunnel is established between the PE router connected to the source site and its remote PE router neighbors.
- The source PE advertises the new data MDT group as long as the source is active. The periodic announcement is sent over the default MDT for the VRF. Because the data MDT announcement is sent over the default tunnel, all the PE routers receive the announcement.
- Neighbors that do not have receivers for the multicast traffic cache the advertisement of the new data MDT group but ignore the new tunnel. Neighbors that do have receivers for the multicast traffic cache the advertisement of the new data MDT group and also send a PIM join message for the new group.
- The source PE encapsulates the VRF multicast traffic using the new data MDT group and stops the packet flow over the default multicast tree. If the multicast traffic level drops back below the threshold, the data MDT is torn down automatically and traffic flows back across the default multicast tree.
- If a PE router that has not yet joined the new data MDT group receives a PIM join message for a new receiver for which (S,G) traffic is already flowing over the data MDT in the provider core, then that PE router can obtain the new group address from its cache and can join the data-MDT immediately without waiting up to 59 seconds for the next data MDT advertisement.

By default, automatic creation of data MDTs is disabled.

The following sections summarize the data MDT configuration statements used in this example and in the prerequisite configuration for this example:

- In the master instance, the PE router's prerequisite draft-rosen PIM-SSM multicast configuration includes statements that directly support the data MDT configuration you will enable in this example. [Table 18 on page 671](#) highlights some of these statements[†].

Table 18: Data MDTs—Key Prerequisites in the Master Instance

Statement	Description
<pre>[edit protocols] pim { interface (Protocols PIM) interface-name <options>; }</pre>	<p>Enables the PIM protocol on PE router interfaces.</p>
<pre>[edit protocols] bgp { group name { type internal; peer-as autonomous-system; neighbor address; family inet-mdt { signaling; } } } [edit routing-options] autonomous-system autonomous-system;</pre>	<p>In the internal BGP full mesh between PE routers in the VRF instance, enables the BGP protocol to carry MDT-SAFI NLRI signaling messages for IPv4 traffic in Layer 3 VPNs.</p>
<pre>[edit routing-options] multicast { ssm-groups [ip-addresses]; }</pre>	<p>(Optional) Configures one or more SSM groups to use inside the provider network in addition to the default SSM group address range of 232.0.0.0/8.</p> <p>NOTE: For this example, it is assumed that you previously specified an additional SSM group address range of 239.0.0.0/8.</p>

† This table contains only a partial list of the PE router configuration statements for a draft-rosen multicast VPN operating in SSM mode in the provider core. For complete configuration information about this prerequisite, see [“Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs” on page 649](#) in the Multicast Protocols User Guide.

- In the VRF instance to which the PE router is attached—at the `[edit routing-instances name]` hierarchy level—the PE router’s prerequisite draft-rosen PIM-SSM multicast configuration includes statements that directly support the data MDT configuration you will enable in this example. [Table 19 on page 672](#) highlights some of these statements[‡].

Table 19: Data MDTs—Key Prerequisites in the VRF Instance

Statement	Description
<pre>[edit routing-instances name] instance-type vrf; vrf-target community;</pre>	<p>Creates a VRF table that contains the routes destined for the PE router.</p> <p>Creates a VRF export list that accepts routes from the routing table. enables using the inet-mcast family.</p> <p>You must also configure the <code>route-distinguish</code> parameter in the routing instance.</p>
<pre>[edit routing-instances name] protocols { pim { mvpn { family { inet / inet6 { autodiscovery { inet-mdt; } } } } } }</pre>	<p>Configures the PIM-SSM MDT-SAFI NLRI for PE routers:</p>

Table 19: Data MDTs—Key Prerequisites in the VRF Instance (*Continued*)

Statement	Description
<pre>[edit routing-instances name] provider-tunnel family inet inet6 { pim-ssm { group-address (Routing Instances) address; } }</pre>	<p>Configures the PIM-SSM provider tunnel for the default MDT group.</p> <p>NOTE: For this configuration to work, you must have previously configured a provider tunnel for the VRF instance ce1 with the address 239.1.1.1.</p> <p>To verify the configuration of the provider tunnel for the VRF instance, use the <code>show</code> command on the router is attached to the VRF instance in operational mode.</p>

‡ This table contains only a partial list of the PE router configuration statements for a draft-rosen multicast VPN operating in SSM mode. For complete configuration information about this prerequisite, see [“Example: Configuring Source-Specific Multicast for Draft-Rosen 7 MVPN”](#) in the Multicast Protocols User Guide.

- For a rosen 7 MVPN—a draft-rosen multicast VPN with provider tunnels operating in SSM mode—you configure data MDT creation for a tunnel multicast group by including statements under the PIM-SSM provider tunnel configuration for the VRF instance associated with the multicast group. Because data MDTs are specific to VPNs and VRF routing instances, you cannot configure MDT statements in the primary routing instance. [Table 20 on page 674](#) summarizes the data MDT configuration statements for PIM-SSM provider tunnels.

Table 20: Data MDTs for PIM-SSM Provider Tunnels in a Draft-Rosen MVPN

Statement	Description
<pre>[edit routing-instances name] provider-tunnel family inet / inet6{ mdt { group-range multicast-prefix; } }</pre>	<p>Configures the IP group range used when a new data MDT needs to be created in the VRF instance on the PE router. This address range cannot overlap the default MDT addresses of any other VPNs on the router. If you configure overlapping group ranges, the configuration commit fails.</p> <p>This statement has no default value. If you do not set the <i>multicast-prefix</i> to a valid, nonreserved multicast address range, then no data MDTs are created for this VRF instance.</p> <p>NOTE: For this example, it is assumed that you previously configured the PE router to automatically select an address from the 239.10.10.0/24 range when a new data MDT needs to be initiated.</p>
<pre>[edit routing-instances name] provider-tunnel family inet / inet6{ mdt { tunnel-limit limit; } }</pre>	<p>Configures the maximum number of data MDTs that can be created for the VRF instance.</p> <p>The default value is 0. If you do not configure the <i>limit</i> to a non-zero value, then no data MDTs are created for this VRF instance.</p> <p>The valid range is from 0 through 1024 for a VRF instance. There is a limit of 8000 tunnels for all data MDTs in all VRF instances on a PE router.</p> <p>If the configured maximum number of data MDT tunnels is reached, then no new tunnels are created for the VRF instance, and traffic that exceeds the configured threshold is sent on the default MDT.</p> <p>NOTE: For this example, you limit the number of data MDTs for the VRF instance to 10.</p>

Table 20: Data MDTs for PIM-SSM Provider Tunnels in a Draft-Rosen MVPN (*Continued*)

Statement	Description
<pre>[edit routing-instances name] provider-tunnel family inet inet6{{ mdt { threshold { group group-address { source source-address { rate threshold-rate; } } } } }</pre>	<p>Configures a data rate for the multicast source of a default MDT. When the source traffic in the VRF instance exceeds the configured data rate, a new tunnel is created.</p> <ul style="list-style-type: none"> • group <i>group-address</i>—Multicast group address of the default MDT that corresponds to a VRF instance to which the PE router is attached. The <i>group-address</i> explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). This is typically a well-known address for a certain type of multicast traffic. • source <i>source-address</i>—Unicast IP prefix of one or more multicast sources in the specified default MDT group. • rate <i>threshold-rate</i>—Data rate for the multicast source to trigger the automatic creation of a data MDT. The data rate is specified in kilobits per second (Kbps). The default <i>threshold-rate</i> is 10 kilobits per second (Kbps). <p>NOTE: For this example, you configure the following data MDT threshold:</p> <ul style="list-style-type: none"> • Multicast group address or address range to which the threshold limits apply—224.0.9.0/32 • Multicast source address or address range to which the threshold limits apply—10.1.1.2/32 • Data rate—10 Kbps <p>When the traffic stops or the rate falls below the threshold value, the source PE router switches back to the default MDT.</p>

Topology

Figure 89 on page 676 shows a default MDT.

Figure 89: Default MDT





Figure 90 on page 676 shows a data MDT.

Figure 90: Data MDT



Configuration

IN THIS SECTION

- CLI Quick Configuration | 677
- Enabling Data MDTs and PIM-SSM Provider Tunnels on the Local PE Router Attached to a VRF | 677
- (Optional) Enabling Logging of Detailed Trace Information for Multicast Tunnel Interfaces on the Local PE Router | 679

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the [Junos OS CLI User Guide](#).

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.

```
set routing-instances ce1 provider-tunnel family inet mdt group-range 239.10.10.0/24
set routing-instances ce1 provider-tunnel family inet mdt tunnel-limit 10
set routing-instances ce1 provider-tunnel family inet mdt threshold group 224.0.9.0/32 source
10.1.1.2/32 rate 10
set protocols pim traceoptions file trace-pim-mdt
set protocols pim traceoptions file files 5
set protocols pim traceoptions file size 1m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag mdt detail
```

Enabling Data MDTs and PIM-SSM Provider Tunnels on the Local PE Router Attached to a VRF

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 [Junos OS CLI User Guide](#).

To configure the local PE router attached to the VRF instance **ce1** in a PIM-SSM multicast VPN to initiate new data MDTs and provider tunnels for that VRF:

1. Enable configuration of provider tunnels operating in SSM mode.

```
[edit]
user@host# edit routing-instances ce1 provider-tunnel
```

2. Configure the range of multicast IP addresses for new data MDTs.

```
[edit routing-instances ce1 provider-tunnel]
user@host# set mdt group-range 239.10.10.0/24
```

3. Configure the maximum number of data MDTs for this VRF instance.

```
[edit routing-instances ce1 provider-tunnel]
user@host# set mdt tunnel-limit 10
```

4. Configure the data MDT-creation threshold for a multicast group and source.

```
[edit routing-instances ce1 provider-tunnel]
user@host# set mdt threshold group 224.0.9.0/32 source 10.1.1.2/32 rate 10
```

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

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

Results

Confirm the configuration of data MDTs for PIM-SSM provider tunnels by entering the `show routing-instances` command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show routing-instances
ce1 {
  instance-type vrf;
  vrf-target target:100:1;
  ...
  provider-tunnel {
    pim-ssm {
      group-address 239.1.1.1;
    }
    mdt {
```



```

        threshold {
            group 224.0.9.0/32 {
                source 10.1.1.2/32 {
                    rate 10;
                }
            }
        }
        tunnel-limit 10;
        group-range 239.10.10.0/24;
    }
}
protocols {
    ...
    pim {
        mvpn {
            family {
                inet {
                    autodiscovery {
                        inet-mdt;
                    }
                }
            }
        }
    }
}
}
}
}
}

```



NOTE: The `show routing-instances` command output above does not show the complete configuration of a VRF instance in a draft-rosen MVPN operating in SSM mode in the provider core.

(Optional) Enabling Logging of Detailed Trace Information for Multicast Tunnel Interfaces on the Local PE Router

Step-by-Step Procedure

To enable logging of detailed trace information for all multicast tunnel interfaces on the local PE router:

1. Enable configuration of PIM tracing options.

```
[edit]
user@host# set protocols pim traceoptions
```

2. Configure the trace file name, maximum number of trace files, maximum size of each trace file, and file access type.

```
[edit protocols pim traceoptions]
set file trace-pim-mdt
set file files 5
set file size 1m
set file world-readable
```

3. Specify that messages related to multicast data tunnel operations are logged.

```
[edit protocols pim traceoptions]
set flag mdt detail
```

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

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

Results

Confirm the configuration of multicast tunnel logging by entering the `show protocols` command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show protocols
pim {
  traceoptions {
    file trace-pim-mdt size 1m files 5 world-readable;
    flag mdt detail;
  }
  interface lo0.0;
```

```
...
}
```

Verification

IN THIS SECTION

- [Monitor Data MDTs Initiated for the Multicast Group | 681](#)
- [Monitor Data MDT Group Addresses Cached by All PE Routers in the Multicast Group | 681](#)
- [\(Optional\) View the Trace Log for Multicast Tunnel Interfaces | 682](#)

To verify that the local PE router is managing data MDTs and PIM-SSM provider tunnels properly, perform the following tasks:

Monitor Data MDTs Initiated for the Multicast Group

Purpose

For the VRF instance **ce1**, check the incoming and outgoing tunnels established by the local PE router for the default MDT and monitor the data MDTs initiated by the local PE router.

Action

Use the **show pim mdt instance ce1 detail** operational mode command.

For the default MDT, the command displays details about the incoming and outgoing tunnels established by the local PE router for specific multicast source addresses in the multicast group using the default MDT and identifies the tunnel mode as **PIM-SSM**.

For the data MDTs initiated by the local PE router, the command identifies the multicast source using the data MDT, the multicast tunnel logical interface set up for the data MDT tunnel, the configured threshold rate, and current statistics.

Monitor Data MDT Group Addresses Cached by All PE Routers in the Multicast Group

Purpose

For the VRF instance **ce1**, check the data MDT group addresses cached by all PE routers that participate in the VRF.

Action

Use the **show pim mdt data-mdt-joins instance ce1** operational mode command. The command output displays the information cached from MDT join TLV packets received by all PE routers participating in the specified VRF instance, including the current timeout value of each entry.

(Optional) View the Trace Log for Multicast Tunnel Interfaces

Purpose

If you configured logging of trace Information for multicast tunnel interfaces, you can trace the creation and tear-down of data MDTs on the local router through the **mt** interface-related activity in the log.

Action

To view the trace file, use the **file show /var/log/trace-pim-mdt** operational mode command.

RELATED DOCUMENTATION

[Load Balancing Multicast Tunnel Interfaces Among Available PICs | 625](#)

[Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs | 649](#)

[Tunnel Services PICs and Multicast | 315](#)

Examples: Configuring Data MDTs

IN THIS SECTION

- [Understanding Data MDTs | 683](#)
- [Data MDT Characteristics | 684](#)
- [Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode | 685](#)
- [Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode | 698](#)
- [Example: Enabling Dynamic Reuse of Data MDT Group Addresses | 704](#)

Understanding Data MDTs

In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, the VPN is multicast-enabled and configured to use the Protocol Independent Multicast (PIM) protocol within the VPN and within the service provider (SP) network. A multicast-enabled VPN routing and forwarding (VRF) instance corresponds to a multicast domain (MD), and a PE router attached to a particular VRF instance is said to belong to the corresponding MD. For each MD there is a default multicast distribution tree (MDT) through the SP backbone, which connects all of the PE routers belonging to that MD. Any PE router configured with a default MDT group address can be the multicast source of one default MDT.

To provide optimal multicast routing, you can configure the PE routers so that when the multicast source within a site exceeds a traffic rate threshold, the PE router to which the source site is attached creates a new data MDT and advertises the new MDT group address. An advertisement of a new MDT group address is sent in a User Datagram Protocol (UDP) type-length-value (TLV) packet called an *MDT join TLV*. The MDT join TLV identifies the source and group pair (S,G) in the VRF instance as well as the new data MDT group address used in the provider space. The PE router to which the source site is attached sends the MDT join TLV over the default MDT for that VRF instance every 60 seconds as long as the source is active.

All PE routers in the VRF instance receive the MDT join TLV because it is sent over the default MDT, but not all the PE routers join the new data MDT group:

- PE routers connected to receivers in the VRF instance for the current multicast group cache the contents of the MDT join TLV, adding a 180-second timeout value to the cache entry, and also join the new data MDT group.
- PE routers not connected to receivers listed in the VRF instance for the current multicast group also cache the contents of the MDT join TLV, adding a 180-second timeout value to the cache entry, but do not join the new data MDT group at this time.

After the source PE stops sending the multicast traffic stream over the default MDT and uses the new MDT instead, only the PE routers that join the new group receive the multicast traffic for that group.

When a remote PE router joins the new data MDT group, it sends a PIM join message for the new group directly to the source PE router from the remote PE routers by means of a PIM (S,G) join.

If a PE router that has not yet joined the new data MDT group receives a PIM join message for a new receiver for which (S,G) traffic is already flowing over the data MDT in the provider core, then that PE router can obtain the new group address from its cache and can join the data MDT immediately without waiting up to 59 seconds for the next data MDT advertisement.

When the PE router to which the source site is attached sends a subsequent MDT join TLV for the VRF instance over the default MDT, any existing cache entries for that VRF instance are simply refreshed with a timeout value of 180 seconds.

To display the information cached from MDT join TLV packets received by all PE routers in a PIM-enabled VRF instance, use the **show pim mdt data-mdt-joins** *operational mode command*.

The source PE router starts encapsulating the multicast traffic for the VRF instance using the new data MDT group after 3 seconds, allowing time for the remote PE routers to join the new group. The source PE router then halts the flow of multicast packets over the default MDT, and the packet flow for the VRF instance source shifts to the newly created data MDT.

The PE router monitors the traffic rate during its periodic statistics-collection cycles. If the traffic rate drops below the threshold or the source stops sending multicast traffic, the PE router to which the source site is attached stops announcing the MDT join TLVs and switches back to sending on the default MDT for that VRF instance.

SEE ALSO

show pim mdt data-mdt-joins

[CLI Explorer](#)

Data MDT Characteristics

A data multicast distribution tree (MDT) solves the problem of routers flooding unnecessary multicast information to PE routers that have no interested receivers for a particular VPN multicast group.

The default MDT uses multicast tunnel (**mt-**) logical interfaces. Data MDTs also use multicast tunnel logical interfaces. If you administratively disable the physical interface that the multicast tunnel logical interfaces are configured on, the multicast tunnel logical interfaces are moved to a different physical interface that is up. In this case the traffic is sent over the default MDT until new data MDTs are created.

The maximum number of data MDTs for all VPNs on a PE router is 1024, and the maximum number of data MDTs for a VRF instance is 1024. The configuration of a VRF instance can limit the number of MDTs possible. No new MDTs can be created after the 1024 MDT limit is reached in the VRF instance, and all traffic for other sources that exceed the configured limit is sent on the default MDT.

Tear-down of data MDTs depends on the monitoring of the multicast source data rate. This rate is checked once per minute, so if the source data rate falls below the configured value, data MDT deletion can be delayed for up to 1 minute until the next statistics-monitoring collection cycle.

Changes to the configured data MDT limit value do not affect existing tunnels that exceed the new limit. Data MDTs that are already active remain in place until the threshold conditions are no longer met.

In a draft-rosen MVPN in which PE routers are already configured to create data MDTs in response to exceeded multicast source traffic rate thresholds, you can change the group range used for creating data MDTs in a VRF instance. To remove any active data MDTs created using the previous group range, you must restart the PIM routing process. This restart clears all remnants of the former group addresses but disrupts routing and therefore requires a maintenance window for the change.



CAUTION: Never restart any of the software processes unless instructed to do so by a customer support engineer.

Multicast tunnel (**mt**) interfaces created because of exceeded thresholds are not re-created if the routing process crashes. Therefore, graceful restart does not automatically reinstate the data MDT state. However, as soon as the periodic statistics collection reveals that the threshold condition is still exceeded, the tunnels are quickly re-created.

Data MDTs are supported for customer traffic with PIM sparse mode, dense mode, and sparse-dense mode. Note that the provider core does not support PIM dense mode.

Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

IN THIS SECTION

- [Requirements](#) | 685
- [Overview](#) | 686
- [Configuration](#) | 692
- [Verification](#) | 697

This example shows how to configure data multicast distribution trees (MDTs) for a provider edge (PE) router attached to a VPN routing and forwarding (VRF) instance in a draft-rosen Layer 3 multicast VPN operating in source-specific multicast (SSM) mode. The example is based on the Junos OS implementation of RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and on section 7 of the IETF Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP IP VPNs*.

Requirements

Before you begin:

- Make sure that the routing devices support multicast tunnel (**mt**) interfaces.

A tunnel-capable PIC supports a maximum of 512 multicast tunnel interfaces. Both default and data MDTs contribute to this total. The default MDT uses two multicast tunnel interfaces (one for encapsulation and one for de-encapsulation). To enable an M Series or T Series router to support more than 512 multicast tunnel interfaces, another tunnel-capable PIC is required. See [""Tunnel](#)

[Services PICs and Multicast" on page 315](#)" and ["Load Balancing Multicast Tunnel Interfaces Among Available PICs" on page 625](#)" in the Multicast Protocols User Guide.

- Make sure that the PE router has been configured for a draft-rosen Layer 3 multicast VPN operating in SSM mode in the provider core.

In this type of multicast VPN, PE routers discover one another by sending MDT subsequent address family identifier (MDT-SAFI) BGP network layer reachability information (NLRI) advertisements. Key configuration statements for the master instance are highlighted in [Table 21 on page 687](#). Key configuration statements for the VRF instance to which your PE router is attached are highlighted in [Table 22 on page 688](#). For complete configuration details, see ["Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs" on page 649](#)" in the Multicast Protocols User Guide.

Overview

IN THIS SECTION

- [Topology | 692](#)

By using data MDTs in a Layer 3 VPN, you can prevent multicast packets from being flooded unnecessarily to specified provider edge (PE) routers within a VPN group. This option is primarily useful for PE routers in your Layer 3 VPN multicast network that have no receivers for the multicast traffic from a particular source.

- When a PE router that is directly connected to the multicast source (also called the *source PE*) receives Layer 3 VPN multicast traffic that exceeds a configured threshold, a new data MDT tunnel is established between the PE router connected to the source site and its remote PE router neighbors.
- The source PE advertises the new data MDT group as long as the source is active. The periodic announcement is sent over the default MDT for the VRF. Because the data MDT announcement is sent over the default tunnel, all the PE routers receive the announcement.
- Neighbors that do not have receivers for the multicast traffic cache the advertisement of the new data MDT group but ignore the new tunnel. Neighbors that do have receivers for the multicast traffic cache the advertisement of the new data MDT group and also send a PIM join message for the new group.
- The source PE encapsulates the VRF multicast traffic using the new data MDT group and stops the packet flow over the default multicast tree. If the multicast traffic level drops back below the

threshold, the data MDT is torn down automatically and traffic flows back across the default multicast tree.

- If a PE router that has not yet joined the new data MDT group receives a PIM join message for a new receiver for which (S,G) traffic is already flowing over the data MDT in the provider core, then that PE router can obtain the new group address from its cache and can join the data-MDT immediately without waiting up to 59 seconds for the next data MDT advertisement.

By default, automatic creation of data MDTs is disabled.

The following sections summarize the data MDT configuration statements used in this example and in the prerequisite configuration for this example:

- In the master instance, the PE router's prerequisite draft-rosen PIM-SSM multicast configuration includes statements that directly support the data MDT configuration you will enable in this example. [Table 21 on page 687](#) highlights some of these statements[†].

Table 21: Data MDTs—Key Prerequisites in the Master Instance

Statement	Description
<pre>[edit protocols] pim { interface (Protocols PIM) interface-name <options>; }</pre>	Enables the PIM protocol on PE router interfaces.
<pre>[edit protocols] bgp { group name { type internal; peer-as autonomous-system; neighbor address; family inet-mdt { signaling; } } } [edit routing-options] autonomous-system autonomous-system;</pre>	In the internal BGP full mesh between PE routers in the VRF instance, enables the BGP protocol to carry MDT-SAFI NLRI signaling messages for IPv4 traffic in Layer 3 VPNs.

Table 21: Data MDTs—Key Prerequisites in the Master Instance (*Continued*)

Statement	Description
<pre>[edit routing-options] multicast { ssm-groups [ip-addresses]; }</pre>	<p>(Optional) Configures one or more SSM groups to use inside the provider network in addition to the default SSM group address range of 232.0.0.0/8.</p> <p>NOTE: For this example, it is assumed that you previously specified an additional SSM group address range of 239.0.0.0/8.</p>

† This table contains only a partial list of the PE router configuration statements for a draft-rosen multicast VPN operating in SSM mode in the provider core. For complete configuration information about this prerequisite, see [“Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs” on page 649](#) in the Multicast Protocols User Guide.

- In the VRF instance to which the PE router is attached—at the `[edit routing-instances name]` hierarchy level—the PE router’s prerequisite draft-rosen PIM-SSM multicast configuration includes statements that directly support the data MDT configuration you will enable in this example. [Table 22 on page 688](#) highlights some of these statements†.

Table 22: Data MDTs—Key Prerequisites in the VRF Instance

Statement	Description
<pre>[edit routing-instances name] instance-type vrf; vrf-target community;</pre>	<p>Creates a VRF table that contains the routes destined for the VRF.</p> <p>Creates a VRF export list that accepts routes from the VRF routing table. Ensures that routes are exported using the inet-mcast export policy.</p> <p>You must also configure a route-distinguisher for the VRF routing instance.</p>

Table 22: Data MDTs—Key Prerequisites in the VRF Instance (*Continued*)

Statement	Description
<pre>[edit routing-instances name] protocols { pim { mvpn { family { inet inet6 { autodiscovery { inet-mdt; } } } } } }</pre>	<p>Configures the PIM MDT-SAFI NLRI for all routers:</p>
<pre>[edit routing-instances name] provider-tunnel family inet inet6 { pim-ssm { group-address (Routing Instances) address; } }</pre>	<p>Configures the PIM default MDT group.</p> <p>NOTE: For this configuration, you previously configured a provider tunnel instance ce1 with address 239.1.1.1.</p> <p>To verify the configuration of the provider tunnel for the VRF instance, use the <code>show</code> command on the router is attached to the VRF instance in operational mode.</p>

‡ This table contains only a partial list of the PE router configuration statements for a draft-rosen multicast VPN operating in SSM mode. For complete configuration information about this prerequisite, see [“Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPN 649”](#) in the Multicast Protocols User Guide.

- For a rosen 7 MVPN—a draft-rosen multicast VPN with provider tunnels operating in SSM mode—you configure data MDT creation for a tunnel multicast group by including statements under the PIM-SSM provider tunnel configuration for the VRF instance associated with the multicast group. Because data MDTs are specific to VPNs and VRF routing instances, you cannot configure MDT

statements in the primary routing instance. [Table 23 on page 690](#) summarizes the data MDT configuration statements for PIM-SSM provider tunnels.

Table 23: Data MDTs for PIM-SSM Provider Tunnels in a Draft-Rosen MVPN

Statement	Description
<pre>[edit routing-instances name] provider-tunnel family inet inet6{{ mdt { group-range multicast-prefix; } }}</pre>	<p>Configures the IP group range used when a new data MDT needs to be created in the VRF instance on the PE router. This address range cannot overlap the default MDT addresses of any other VPNs on the router. If you configure overlapping group ranges, the configuration commit fails.</p> <p>This statement has no default value. If you do not set the <i>multicast-prefix</i> to a valid, nonreserved multicast address range, then no data MDTs are created for this VRF instance.</p> <p>NOTE: For this example, it is assumed that you previously configured the PE router to automatically select an address from the 239.10.10.0/24 range when a new data MDT needs to be initiated.</p>
<pre>[edit routing-instances name] provider-tunnel family inet inet6{{ mdt { tunnel-limit limit; } }}</pre>	<p>Configures the maximum number of data MDTs that can be created for the VRF instance.</p> <p>The default value is 0. If you do not configure the <i>limit</i> to a non-zero value, then no data MDTs are created for this VRF instance.</p> <p>The valid range is from 0 through 1024 for a VRF instance. There is a limit of 8000 tunnels for all data MDTs in all VRF instances on a PE router.</p> <p>If the configured maximum number of data MDT tunnels is reached, then no new tunnels are created for the VRF instance, and traffic that exceeds the configured threshold is sent on the default MDT.</p> <p>NOTE: For this example, you limit the number of data MDTs for the VRF instance to 10.</p>

Table 23: Data MDTs for PIM-SSM Provider Tunnels in a Draft-Rosen MVPN (Continued)

Statement	Description
<pre>[edit routing-instances name] provider-tunnel family inet / inet6{ mdt { threshold { group group-address { source source-address { rate threshold-rate; } } } } }</pre>	<p>Configures a data rate for the multicast source of a default MDT. When the source traffic in the VRF instance exceeds the configured data rate, a new tunnel is created.</p> <ul style="list-style-type: none"> • group <i>group-address</i>—Multicast group address of the default MDT that corresponds to a VRF instance to which the PE router is attached. The <i>group-address</i> explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). This is typically a well-known address for a certain type of multicast traffic. • source <i>source-address</i>—Unicast IP prefix of one or more multicast sources in the specified default MDT group. • rate <i>threshold-rate</i>—Data rate for the multicast source to trigger the automatic creation of a data MDT. The data rate is specified in kilobits per second (Kbps). The default <i>threshold-rate</i> is 10 kilobits per second (Kbps). <p>NOTE: For this example, you configure the following data MDT threshold:</p> <ul style="list-style-type: none"> • Multicast group address or address range to which the threshold limits apply—224.0.9.0/32 • Multicast source address or address range to which the threshold limits apply—10.1.1.2/32 • Data rate—10 Kbps <p>When the traffic stops or the rate falls below the threshold value, the source PE router switches back to the default MDT.</p>

Topology

Figure 91 on page 692 shows a default MDT.

Figure 91: Default MDT

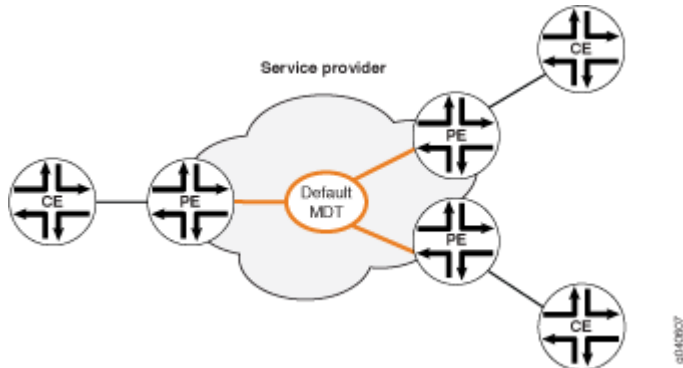



Figure 92 on page 692 shows a data MDT.

Figure 92: Data MDT



Configuration

IN THIS SECTION

- CLI Quick Configuration | 693
- Enabling Data MDTs and PIM-SSM Provider Tunnels on the Local PE Router Attached to a VRF | 693
- (Optional) Enabling Logging of Detailed Trace Information for Multicast Tunnel Interfaces on the Local PE Router | 695

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the [Junos OS CLI User Guide](#).

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.

```
set routing-instances ce1 provider-tunnel family inet mdt group-range 239.10.10.0/24
set routing-instances ce1 provider-tunnel family inet mdt tunnel-limit 10
set routing-instances ce1 provider-tunnel family inet mdt threshold group 224.0.9.0/32 source
10.1.1.2/32 rate 10
set protocols pim traceoptions file trace-pim-mdt
set protocols pim traceoptions file files 5
set protocols pim traceoptions file size 1m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag mdt detail
```

Enabling Data MDTs and PIM-SSM Provider Tunnels on the Local PE Router Attached to a VRF

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 [Junos OS CLI User Guide](#).

To configure the local PE router attached to the VRF instance **ce1** in a PIM-SSM multicast VPN to initiate new data MDTs and provider tunnels for that VRF:

1. Enable configuration of provider tunnels operating in SSM mode.

```
[edit]
user@host# edit routing-instances ce1 provider-tunnel
```

2. Configure the range of multicast IP addresses for new data MDTs.

```
[edit routing-instances ce1 provider-tunnel]
user@host# set mdt group-range 239.10.10.0/24
```

3. Configure the maximum number of data MDTs for this VRF instance.

```
[edit routing-instances ce1 provider-tunnel]
user@host# set mdt tunnel-limit 10
```

4. Configure the data MDT-creation threshold for a multicast group and source.

```
[edit routing-instances ce1 provider-tunnel]
user@host# set mdt threshold group 224.0.9.0/32 source 10.1.1.2/32 rate 10
```

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

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

Results

Confirm the configuration of data MDTs for PIM-SSM provider tunnels by entering the `show routing-instances` command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show routing-instances
ce1 {
  instance-type vrf;
  vrf-target target:100:1;
  ...
  provider-tunnel {
    pim-ssm {
      group-address 239.1.1.1;
    }
    mdt {
```



```

threshold {
    group 224.0.9.0/32 {
        source 10.1.1.2/32 {
            rate 10;
        }
    }
}
tunnel-limit 10;
group-range 239.10.10.0/24;
}
}
protocols {
    ...
    pim {
        mvpn {
            family {
                inet {
                    autodiscovery {
                        inet-mdt;
                    }
                }
            }
        }
    }
}
}
}
}
}

```



NOTE: The `show routing-instances` command output above does not show the complete configuration of a VRF instance in a draft-rosen MVPN operating in SSM mode in the provider core.

(Optional) Enabling Logging of Detailed Trace Information for Multicast Tunnel Interfaces on the Local PE Router

Step-by-Step Procedure

To enable logging of detailed trace information for all multicast tunnel interfaces on the local PE router:

1. Enable configuration of PIM tracing options.

```
[edit]
user@host# set protocols pim traceoptions
```

2. Configure the trace file name, maximum number of trace files, maximum size of each trace file, and file access type.

```
[edit protocols pim traceoptions]
set file trace-pim-mdt
set file files 5
set file size 1m
set file world-readable
```

3. Specify that messages related to multicast data tunnel operations are logged.

```
[edit protocols pim traceoptions]
set flag mdt detail
```

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

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

Results

Confirm the configuration of multicast tunnel logging by entering the `show protocols` command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show protocols
pim {
  traceoptions {
    file trace-pim-mdt size 1m files 5 world-readable;
    flag mdt detail;
  }
  interface lo0.0;
```

```
...
}
```

Verification

IN THIS SECTION

- [Monitor Data MDTs Initiated for the Multicast Group | 697](#)
- [Monitor Data MDT Group Addresses Cached by All PE Routers in the Multicast Group | 698](#)
- [\(Optional\) View the Trace Log for Multicast Tunnel Interfaces | 698](#)

To verify that the local PE router is managing data MDTs and PIM-SSM provider tunnels properly, perform the following tasks:

Monitor Data MDTs Initiated for the Multicast Group

Purpose

For the VRF instance **ce1**, check the incoming and outgoing tunnels established by the local PE router for the default MDT and monitor the data MDTs initiated by the local PE router.

Action

Use the **show pim mdt instance ce1 detail** operational mode command.

For the default MDT, the command displays details about the incoming and outgoing tunnels established by the local PE router for specific multicast source addresses in the multicast group using the default MDT and identifies the tunnel mode as **PIM-SSM**.

For the data MDTs initiated by the local PE router, the command identifies the multicast source using the data MDT, the multicast tunnel logical interface set up for the data MDT tunnel, the configured threshold rate, and current statistics.

Monitor Data MDT Group Addresses Cached by All PE Routers in the Multicast Group

Purpose

For the VRF instance **ce1**, check the data MDT group addresses cached by all PE routers that participate in the VRF.

Action

Use the **show pim mdt data-mdt-joins instance ce1** operational mode command. The command output displays the information cached from MDT join TLV packets received by all PE routers participating in the specified VRF instance, including the current timeout value of each entry.

(Optional) View the Trace Log for Multicast Tunnel Interfaces

Purpose

If you configured logging of trace information for multicast tunnel interfaces, you can trace the creation and tear-down of data MDTs on the local router through the **mt** interface-related activity in the log.

Action

To view the trace file, use the **file show /var/log/trace-pim-mdt** operational mode command.

SEE ALSO

[Load Balancing Multicast Tunnel Interfaces Among Available PICs | 625](#)

[Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs | 649](#)

[Tunnel Services PICs and Multicast | 315](#)

Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

IN THIS SECTION

● [Requirements | 699](#)

● [Overview | 699](#)

● [Configuration | 702](#)

● Verification | 703

This example shows how to configure data multicast distribution trees (MDTs) in a draft-rosen Layer 3 VPN operating in any-source multicast (ASM) mode. This example is based on the Junos OS implementation of RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)* and on section 2 of the IETF Internet draft draft-rosen-vpn-mcast-06.txt, *Multicast in MPLS/BGP VPNs* (expired April 2004).

Requirements

Before you begin:

- Configure the draft-rosen multicast over Layer 3 VPN scenario.
- Make sure that the routing devices support multicast tunnel (**mt**) interfaces.

A tunnel-capable PIC supports a maximum of 512 multicast tunnel interfaces. Both default and data MDTs contribute to this total. The default MDT uses two multicast tunnel interfaces (one for encapsulation and one for de-encapsulation). To enable an M Series or T Series router to support more than 512 multicast tunnel interfaces, another tunnel-capable PIC is required. See ["Tunnel Services PICs and Multicast" on page 315](#) and ["Load Balancing Multicast Tunnel Interfaces Among Available PICs" on page 625](#) in the Multicast Protocols User Guide..

Overview

IN THIS SECTION

● Topology | 701

By using data multicast distribution trees (MDTs) in a Layer 3 VPN, you can prevent multicast packets from being flooded unnecessarily to specified provider edge (PE) routers within a VPN group. This option is primarily useful for PE routers in your Layer 3 VPN multicast network that have no receivers for the multicast traffic from a particular source.

When a PE router that is directly connected to the multicast source (also called the *source PE*) receives Layer 3 VPN multicast traffic that exceeds a configured threshold, a new data MDT tunnel is established between the PE router connected to the source site and its remote PE router neighbors.

The source PE advertises the new data MDT group as long as the source is active. The periodic announcement is sent over the default MDT for the VRF. Because the data MDT announcement is sent over the default tunnel, all the PE routers receive the announcement.

Neighbors that do not have receivers for the multicast traffic cache the advertisement of the new data MDT group but ignore the new tunnel. Neighbors that do have receivers for the multicast traffic cache the advertisement of the new data MDT group and also send a PIM join message for the new group.

The source PE encapsulates the VRF multicast traffic using the new data MDT group and stops the packet flow over the default multicast tree. If the multicast traffic level drops back below the threshold, the data MDT is torn down automatically and traffic flows back across the default multicast tree.

If a PE router that has not yet joined the new data MDT group receives a PIM join message for a new receiver for which (S,G) traffic is already flowing over the data MDT in the provider core, then that PE router can obtain the new group address from its cache and can join the data-MDT immediately without waiting up to 59 seconds for the next data MDT advertisement.

By default, automatic creation of data MDTs is disabled.

For a rosen 6 MVPN—a draft-rosen multicast VPN with provider tunnels operating in ASM mode—you configure data MDT creation for a tunnel multicast group by including statements under the PIM protocol configuration for the VRF instance associated with the multicast group. Because data MDTs apply to VPNs and VRF routing instances, you cannot configure MDT statements in the master routing instance.

This example includes the following configuration options:

- **group**—Specifies the multicast group address to which the threshold applies. This could be a well-known address for a certain type of multicast traffic.

The group address can be explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). Explicit and prefix address forms can be combined if they do not overlap. Overlapping configurations, in which prefix and more explicit address forms are used for the same source or group address, are not supported.

- **group-range**—Specifies the multicast group IP address range used when a new data MDT needs to be initiated on the PE router. For each new data MDT, one address is automatically selected from the configured group range.

The PE router implementing data MDTs for a local multicast source must be configured with a range of multicast group addresses. Group addresses that fall within the configured range are used in the join messages for the data MDTs created in this VRF instance. Any multicast address range can be used as the multicast prefix. However, the group address range cannot overlap the default MDT group address configured for any VPN on the router. If you configure overlapping group addresses, the configuration commit operation fails.

- **pim**—Supports data MDTs for service provider tunnels operating in any-source multicast mode.

- **rate**—Specifies the data rate that initiates the creation of data MDTs. When the source traffic in the VRF exceeds the configured data rate, a new tunnel is created. The range is from 10 kilobits per second (Kbps), the default, to 1 gigabit per second (Gbps, equivalent to 1,000,000 Kbps).
- **source**—Specifies the unicast address of the source of the multicast traffic. It can be a source locally attached to or reached through the PE router. A group can have more than one source.

The source address can be explicit (all 32 bits of the address specified) or a prefix (network address and prefix length specified). Explicit and prefix address forms can be combined if they do not overlap. Overlapping configurations, in which prefix and more explicit address forms are used for the same source or group address, are not supported.

- **threshold**—Associates a rate with a group and a source. The PE router implementing data MDTs for a local multicast source must establish a data MDT-creation threshold for a multicast group and source.

When the traffic stops or the rate falls below the threshold value, the source PE router switches back to the default MDT.

- **tunnel-limit**—Specifies the maximum number of data MDTs that can be created for a single routing instance. The PE router implementing a data MDT for a local multicast source must establish a limit for the number of data MDTs created in this VRF instance. If the limit is 0 (the default), then no data MDTs are created for this VRF instance.

If the number of data MDT tunnels exceeds the maximum configured tunnel limit for the VRF, then no new tunnels are created. Traffic that exceeds the configured threshold is sent on the default MDT.

The valid range is from 0 through 1024 for a VRF instance. There is a limit of 8000 tunnels for all data MDTs in all VRF instances on a PE router.

Topology

Figure 93 on page 701 shows a default MDT.

Figure 93: Default MDT





Figure 94 on page 702 shows a data MDT.

Figure 94: Data MDT



Configuration

IN THIS SECTION

- [Procedure | 702](#)

Procedure

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.

```
[edit]
set routing-instances vpn-A protocols pim mdt group-range 227.0.0.0/8
set routing-instances vpn-A protocols pim mdt threshold group 224.4.4.4/32 source 10.10.20.43/32
rate 10
set routing-instances vpn-A protocols pim mdt tunnel-limit 10
```


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 [Junos OS CLI User Guide](#).

To configure a PE router attached to the VRF instance **vpn-A** in a PIM-ASM multicast VPN to initiate new data MDTs and provider tunnels for that VRF:

1. Configure the group range.

```
[edit]
user@host# edit routing-instances vpn-A protocols pim mdt
[edit routing-instances vpn-A protocols pim mdt]
user@host# set group-range 227.0.0.0/8
```

2. Configure a data MDT-creation threshold for a multicast group and source.

```
[edit routing-instances vpn-A protocols pim mdt]
user@host# set threshold group 224.4.4.4 source 10.10.20.43 rate 10
```

3. Configure a tunnel limit.

```
[edit routing-instances vpn-A protocols pim mdt]
user@host# set tunnel-limit 10
```

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

```
[edit routing-instances vpn-A protocols pim mdt]
user@host# commit
```

Verification

To display information about the default MDT and any data MDTs for the VRF instance **vpn-A**, use the **show pim mdt instance ce1 detail** operational mode command. This command displays either the outgoing tunnels (the tunnels initiated by the local PE router), the incoming tunnels (tunnels initiated by the remote PE routers), or both.

To display the data MDT group addresses cached by PE routers that participate in the VRF instance **vpn-A**, use the **show pim mdt data-mdt-joins instance vpn-A** operational mode command. The command displays the information cached from MDT join TLV packets received by all PE routers participating in the specified VRF instance.

You can trace the operation of data MDTs by including the **mdt detail** flag in the [edit protocols pim traceoptions] configuration. When this flag is set, all the **mt** interface-related activity is logged in trace files.

SEE ALSO

Introduction to Configuring Layer 3 VPNs

[Junos OS VPNs Library for Routing Devices](#)

Example: Enabling Dynamic Reuse of Data MDT Group Addresses

IN THIS SECTION

- [Requirements | 704](#)
- [Overview | 705](#)
- [Configuration | 706](#)
- [Verification | 714](#)

This example describes how to enable dynamic reuse of data multicast distribution tree (MDT) group addresses.

Requirements

Before you begin:

- Configure the router interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure PIM Sparse Mode on the interfaces. See "[Enabling PIM Sparse Mode](#)" on page 316.

Overview

IN THIS SECTION

- [Topology | 705](#)

A limited number of multicast group addresses are available for use in data MDT tunnels. By default, when the available multicast group addresses are all used, no new data MDTs can be created.

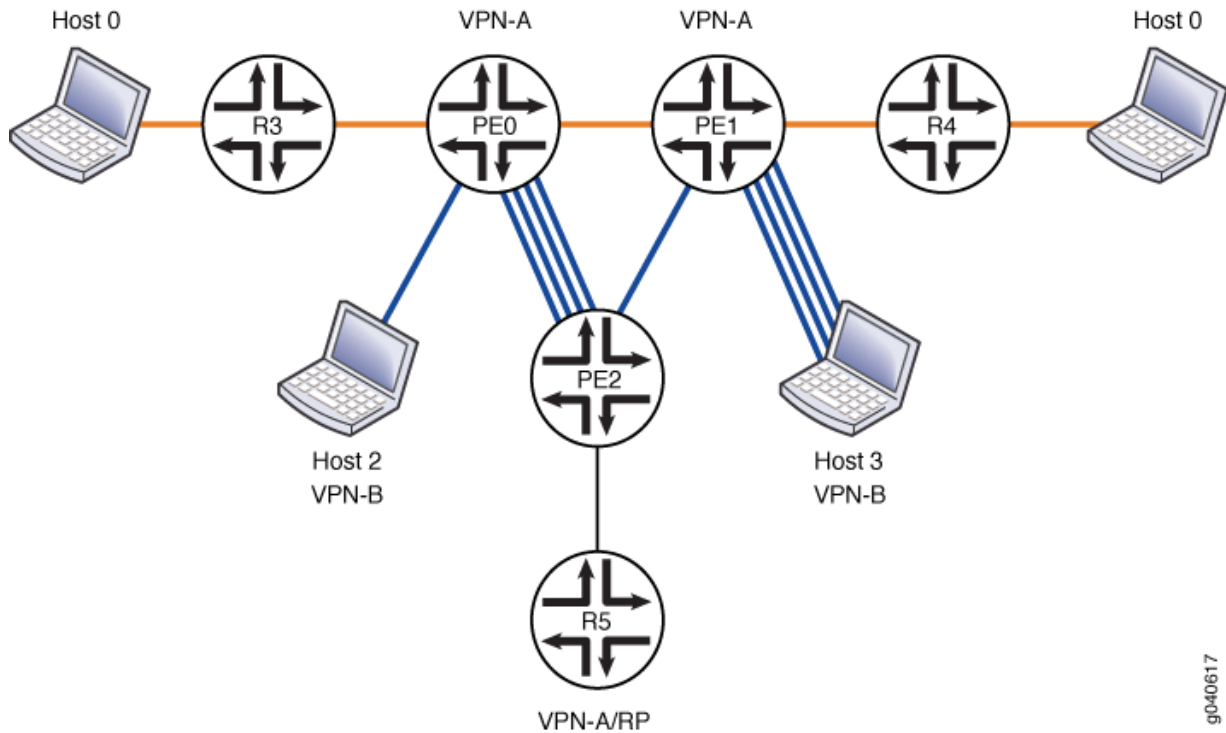
You can enable dynamic reuse of data MDT group addresses. Dynamic reuse of data MDT group addresses allows multiple multicast streams to share a single MDT and multicast provider group address. For example, three streams can use the same provider group address and MDT tunnel.

The streams are assigned to a particular MDT in a round-robin fashion. Since a provider tunnel might be used by multiple customer streams, this can result in egress routers receiving customer traffic that is not destined for their attached customer sites. This example shows the plain PIM scenario, without the MVPN provider tunnel.

Topology

[Figure 95 on page 706](#) shows the topology used in this example.

Figure 95: Dynamic Reuse of Data MDT Group Addresses



9040617

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 706](#)
- [Procedure | 708](#)
- [Results | 711](#)

CLI Quick Configuration

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

```
set 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 protocols mpls interface all
set protocols bgp local-as 65520
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.38.17
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp neighbor 10.255.38.21
set protocols bgp group ibgp neighbor 10.255.38.15
set protocols ospf traffic-engineering
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 pim rp static address 10.255.38.21
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface ge-1/1/2.0
set routing-instances VPN-A interface lo0.1
set routing-instances VPN-A route-distinguisher 10.0.0.10:04
set routing-instances VPN-A vrf-target target:100:10
set routing-instances VPN-A protocols ospf export bgp-to-ospf
set routing-instances VPN-A protocols ospf area 0.0.0.0 interface all
set routing-instances VPN-A protocols pim traceoptions file pim-VPN-A.log
set routing-instances VPN-A protocols pim traceoptions file size 5m
set routing-instances VPN-A protocols pim traceoptions flag mdt detail
set routing-instances VPN-A protocols pim dense-groups 224.0.1.39/32
set routing-instances VPN-A protocols pim dense-groups 224.0.1.40/32
set routing-instances VPN-A protocols pim dense-groups 229.0.0.0/8
set routing-instances VPN-A protocols pim vpn-group-address 239.1.0.0
set routing-instances VPN-A protocols pim rp static address 10.255.38.15
set routing-instances VPN-A protocols pim interface lo0.1 mode sparse-dense
set routing-instances VPN-A protocols pim interface ge-1/1/2.0 mode sparse-dense
set routing-instances VPN-A protocols pim mdt threshold group 224.1.1.1/32 source
192.168.255.245/32 rate 20
set routing-instances VPN-A protocols pim mdt threshold group 224.1.1.2/32 source
192.168.255.245/32 rate 20
set routing-instances VPN-A protocols pim mdt threshold group 224.1.1.3/32 source
192.168.255.245/32 rate 20
set routing-instances VPN-A protocols pim mdt data-mdt-reuse
set routing-instances VPN-A protocols pim mdt tunnel-limit 2
set routing-instances VPN-A protocols pim mdt group-range 239.1.1.0/30
```

Procedure

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 [Junos OS CLI User Guide](#).

To configure dynamic reuse of data MDT group addresses:

1. Configure the **bgp-to-ospf** export policy.

```
[edit policy-options policy-statement bgp-to-ospf]
user@host# set term 1 from protocol bgp
user@host# set term 1 then accept
```

2. Configure MPLS, LDP, BGP, OSPF, and PIM.

```
[edit]
user@host# edit protocols
[edit protocols]
user@host# set mpls interface all
[edit protocols]
user@host# set ldp interface all
[edit protocols]
user@host# set bgp local-as 65520
[edit protocols]
user@host# set bgp group ibgp type internal
[edit protocols]
user@host# set bgp group ibgp local-address 10.255.38.17
[edit protocols]
user@host# set bgp group ibgp family inet-vpn unicast
[edit protocols]
user@host# set bgp group ibgp neighbor 10.255.38.21
[edit protocols]
user@host# set bgp group ibgp neighbor 10.255.38.15
[edit protocols]
user@host# set ospf traffic-engineering
[edit protocols]
user@host# set ospf area 0.0.0.0 interface all
[edit protocols]
user@host# set ospf area 0.0.0.0 interface fxp0.0 disable
```

```

[edit protocols]
user@host# set pim rp static address 10.255.38.21
[edit protocols]
user@host# set pim interface all mode sparse
[edit protocols]
user@host# set pim interface all version 2
[edit protocols]
user@host# set pim interface fxp0.0 disable
[edit protocols]
user@host# exit

```

3. Configure the routing instance, and apply the **bgp-to-ospf** export policy.

```

[edit]
user@host# edit routing-instances VPN-A
[edit routing-instances VPN-A]
user@host# set instance-type vrf
[edit routing-instances VPN-A]
user@host# set interface ge-1/1/2.0
[edit routing-instances VPN-A]
user@host# set interface lo0.1
[edit routing-instances VPN-A]
user@host# set route-distinguisher 10.0.0.10:04
[edit routing-instances VPN-A]
user@host# set vrf-target target:100:10
[edit routing-instances VPN-A]
user@host# set protocols ospf export bgp-to-ospf
[edit routing-instances VPN-A]
user@host# set protocols ospf area 0.0.0.0 interface all

```

4. Configure PIM trace operations for troubleshooting.

```

[edit routing-instances VPN-A]
user@host# set protocols pim traceoptions file pim-VPN-A.log
[edit routing-instances VPN-A]
user@host# set protocols pim traceoptions file size 5m
[edit routing-instances VPN-A]
user@host# set protocols pim traceoptions flag mdt detail

```

5. Configure the groups that operate in dense mode and the group address on which to encapsulate multicast traffic from the routing instance.

```
[edit routing-instances VPN-A]
user@host# set protocols pim dense-groups 224.0.1.39/32
[edit routing-instances VPN-A]
user@host# set protocols pim dense-groups 224.0.1.40/32
[edit routing-instances VPN-A]
user@host# set protocols pim dense-groups 229.0.0.0/8
[edit routing-instances VPN-A]
user@host# set protocols pim group-address 239.1.0.0
[edit routing-instances VPN-A]
```

6. Configure the address of the RP and the interfaces operating in sparse-dense mode.

```
[edit routing-instances VPN-A]
user@host# set protocols pim rp static address 10.255.38.15
[edit routing-instances VPN-A]
user@host# set protocols pim interface lo0.1 mode sparse-dense
[edit routing-instances VPN-A]
user@host# set protocols pim interface ge-1/1/2.0 mode sparse-dense
```

7. Configure the data MDT, including the data-mdt-reuse statement.

```
[edit routing-instances VPN-A]
user@host# set protocols pim mdt threshold group 224.1.1.1/32 source 192.168.255.245/32 rate
20
[edit routing-instances VPN-A]
user@host# set protocols pim mdt threshold group 224.1.1.2/32 source 192.168.255.245/32 rate
20
[edit routing-instances VPN-A]
user@host# set protocols pim mdt threshold group 224.1.1.3/32 source 192.168.255.245/32 rate
20
[edit routing-instances VPN-A]
user@host# set protocols pim mdt data-mdt-reuse
[edit routing-instances VPN-A]
user@host# set protocols pim mdt tunnel-limit 2
[edit routing-instances VPN-A]
user@host# set protocols pim mdt group-range 239.1.1.0/30
```


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

```
[edit routing-instances VPN-A]
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the `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@host# show policy-options
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
}
```

```
user@host# show protocols
mpls {
  interface all;
}
bgp {
  local-as 65520;
  group ibgp {
    type internal;
    local-address 10.255.38.17;
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.38.21;
    neighbor 10.255.38.15;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
```

```
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
}
pim {
    rp {
        static {
            address 10.255.38.21;
        }
    }
    interface all {
        mode sparse;
        version 2;
    }
    interface fxp0.0 {
        disable;
    }
}
```

```
user@host# show routing-instances
VPN-A {
    instance-type vrf;
    interface ge-1/1/2.0;
    interface lo0.1;
    route-distinguisher 10.0.0.10:04;
    vrf-target target:100:10;
    protocols {
        ospf {
            export bgp-to-ospf;
            area 0.0.0.0 {
                interface all;
            }
        }
        pim {
            traceoptions {
                file pim-VPN-A.log size 5m;
                flag mdt detail;
            }
        }
    }
}
```

```
}
dense-groups {
    224.0.1.39/32;
    224.0.1.40/32;
    229.0.0.0/8;
}
vpn-group-address 239.1.0.0;
rp {
    static {
        address 10.255.38.15;
    }
}
interface lo0.1 {
    mode sparse-dense;
}
interface ge-1/1/2.0 {
    mode sparse-dense;
}
mdt {
    threshold {
        group 224.1.1.1/32 {
            source 192.168.255.245/32 {
                rate 20;
            }
        }
        group 224.1.1.2/32 {
            source 192.168.255.245/32 {
                rate 20;
            }
        }
        group 224.1.1.3/32 {
            source 192.168.255.245/32 {
                rate 20;
            }
        }
    }
    data-mdt-reuse;
    tunnel-limit 2;
    group-range 239.1.1.0/30;
}
}
```

```
}  
}
```

Verification

To verify the configuration, run the following commands:

- **show pim join instance VPN-A extensive**
- **show multicast route instance VPN-A extensive**
- **show pim mdt instance VPN-A**
- **show pim mdt data-mdt-joins instance VPN-A**

SEE ALSO

[Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode | 663](#)

RELATED DOCUMENTATION

[Example: Configuring Any-Source Draft-Rosen 6 Multicast VPNs | 611](#)

[Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs | 648](#)

Configuring Next-Generation Multicast VPNs

IN THIS CHAPTER

- Understanding Next-Generation MVPN Network Topology | **716**
- Understanding Next-Generation MVPN Concepts and Terminology | **718**
- Understanding Next-Generation MVPN Control Plane | **720**
- Next-Generation MVPN Data Plane Overview | **730**
- Enabling Next-Generation MVPN Services | **734**
- Generating Next-Generation MVPN VRF Import and Export Policies Overview | **737**
- Multiprotocol BGP MVPNs Overview | **741**
- Configuring Multiprotocol BGP Multicast VPNs | **749**
- BGP-MVPN Inter-AS Option B Overview | **854**
- ACX Support for BGP MVPN | **856**
- Example: Configuring MBGP MVPN Extranets | **858**
- Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs | **913**
- Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs | **914**
- Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels | **928**
- Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels | **935**
- Example: Configuring Sender-Based RPF in a BGP MVPN with MLDP Point-to-Multipoint Provider Tunnels | **972**
- Configuring MBGP MVPN Wildcards | **1008**
- Distributing C-Multicast Routes Overview | **1017**
- Exchanging C-Multicast Routes | **1023**
- Generating Source AS and Route Target Import Communities Overview | **1032**
- Originating Type 1 Intra-AS Autodiscovery Routes Overview | **1033**
- Signaling Provider Tunnels and Data Plane Setup | **1037**
- Anti-spoofing support for MPLS labels in BGP/MPLS IP VPNs (Inter-AS Option B) | **1055**
- BGP-MVPN SD-WAN Overlay | **1056**

Understanding Next-Generation MVPN Network Topology

Layer 3 BGP-MPLS virtual private networks (VPNs) are widely deployed in today's networks worldwide. Multicast applications, such as IPTV, are rapidly gaining popularity as is the number of networks with multiple, media-rich services merging over a shared Multiprotocol Label Switching (MPLS) infrastructure. The demand for delivering multicast service across a BGP-MPLS infrastructure in a scalable and reliable way is also increasing.

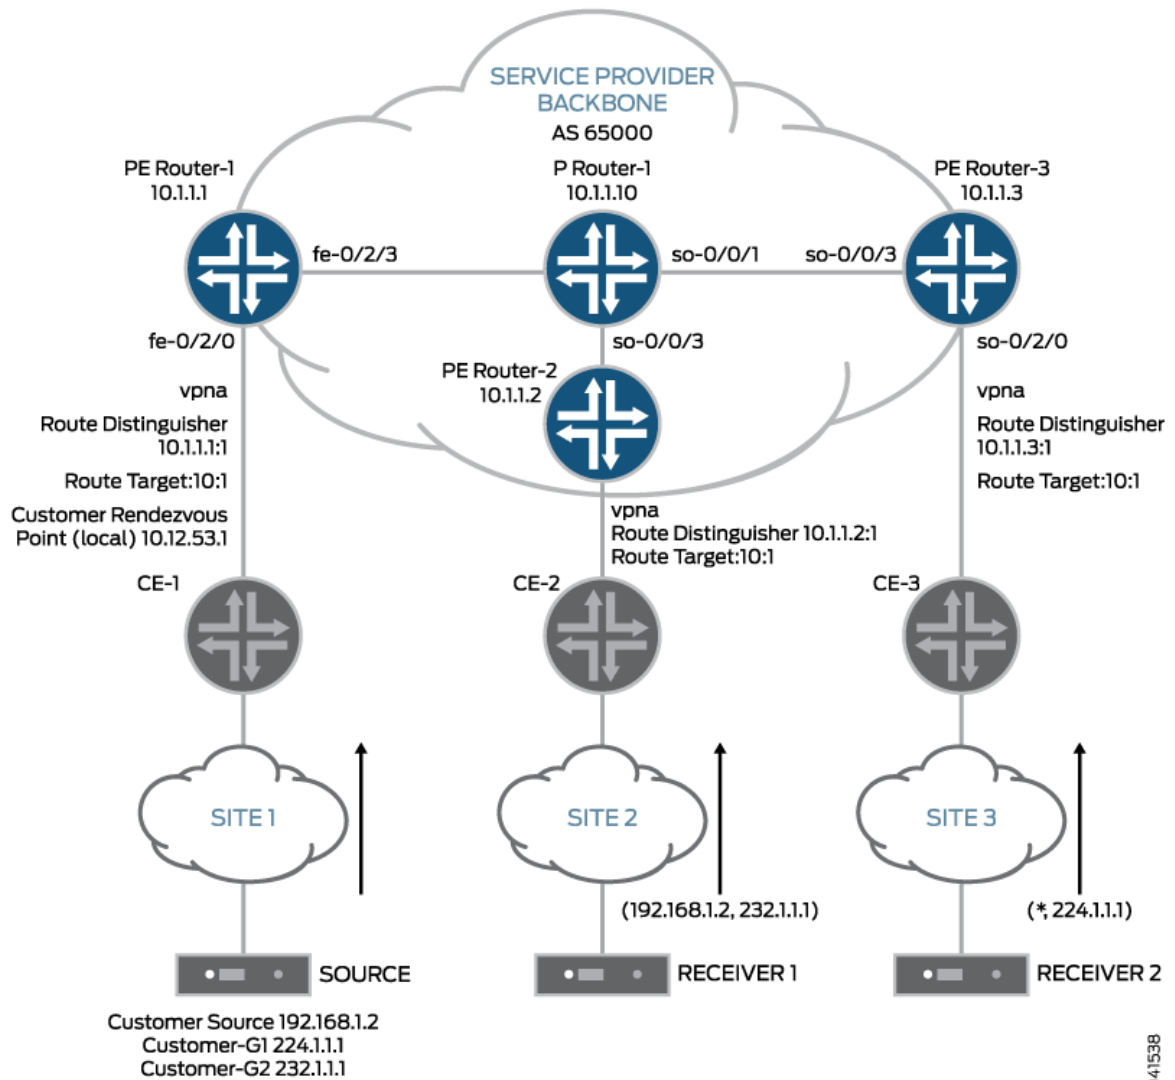
RFC 4364 describes protocols and procedures for building unicast BGP-MPLS VPNs. However, there is no framework specified in the RFC for provisioning multicast VPN (MVPN) services. In the past, Multiprotocol Label Switching Virtual Private Network (MVPN) traffic was overlaid on top of a BGP-MPLS network using a virtual LAN model based on Draft Rosen. Using the Draft Rosen approach, service providers were faced with control and data plane scaling issues of an overlay model and the maintenance of two routing/forwarding mechanisms: one for VPN unicast service and one for VPN multicast service. For more information about the limitations of Draft Rosen, see [draft-rekhter-mboned-mvpn-deploy](#).

As a result, the IETF Layer 3 VPN working group published an Internet draft [draft-ietf-l3vpn-2547bis-mcast-10.txt](#), *Multicast in MPLS/BGP IP VPNs*, that outlines a different architecture for next-generation MVPNs, as well as an accompanying RFC 2547 that proposes a BGP control plane for MVPNs. In turn, Juniper Networks delivered the industry's first implementation of BGP next-generation MVPNs in 2007.

All examples in this document refer to the network topology shown in [Figure 96 on page 717](#):

- The service provider in this example offers VPN unicast and multicast services to Customer A (vpna).
- The VPN multicast source is connected to Site 1 and transmits data to groups 232.1.1.1 and 224.1.1.1.
- VPN multicast receivers are connected to Site 2 and Site 3.
- The provider edge router 1 (Router PE1) VRF table acts as the C-RP (using address 10.12.53.1) for C-PIM-SM ASM groups.
- The service provider uses RSVP-TE point-to-multipoint LSPs for transmitting VPN multicast data across the network.

Figure 96: Next-Generation MVPN Topology



RELATED DOCUMENTATION

[Understanding Next-Generation MVPN Concepts and Terminology](#)

[Understanding Next-Generation MVPN Control Plane | 720](#)

[Next-Generation MVPN Data Plane Overview | 730](#)

Example: Configuring MBGP Multicast VPNs

Understanding Next-Generation MVPN Concepts and Terminology

IN THIS SECTION

- [Route Distinguisher and VRF Route Target Extended Community | 718](#)
- [C-Multicast Routing | 719](#)
- [BGP MVPNs | 719](#)
- [Sender and Receiver Site Sets | 719](#)
- [Provider Tunnels | 720](#)

This section includes background material about how next-generation MVPNs work.

Route Distinguisher and VRF Route Target Extended Community

Route distinguisher and VPN routing and forwarding (VRF) route target extended communities are an integral part of unicast BGP-MPLS virtual private networks (VPNs). Route distinguisher and route target are often confused in terms of their purpose in BGP-MPLS networks. As they play an important role in BGP next-generation MVPNs, it is important to understand what they are and how they are used as described in RFC 4364.

RFC 4364 describes the purpose of route distinguisher as the following:

“A VPN-IPv4 address is a 12-byte quantity, beginning with an 8-byte Route Distinguisher (RD) and ending with a 4-byte IPv4 address. If several VPNs use the same IPv4 address prefix, the PEs translate these into unique VPN-IPv4 address prefixes. This ensures that if the same address is used in several different VPNs, it is possible for BGP to carry several completely different routes to that address, one for each VPN.”

Typically, each VRF table on a provider edge (PE) router is configured with a unique route distinguisher. Depending on the routing design, the route distinguisher can be unique or the same for a given VRF on other PE routers. A route distinguisher is an 8-byte number with two fields. The first field can be either an AS number (2 or 4 bytes) or an IP address (4 bytes). The second field is assigned by the user.

RFC 4364 describes the purpose of a VRF route target extended community as the following:

“Every VRF is associated with one or more Route Target (RT) attributes.”

When a VPN-IPv4 route is created (from an IPv4 route that the PE router has learned from a CE) by a PE router, it is associated with one or more route target attributes. These are carried in BGP as attributes of the route.

Any route associated with Route Target T must be distributed to every PE router that has a VRF associated with Route Target T. When such a route is received by a PE router, it is eligible to be installed in those of the PE's VRFs that are associated with Route Target T."

The route target also contains two fields and is structured similar to a route distinguisher. The first field of the route target is either an AS number (2 or 4 bytes) or an IP address (4 bytes), and the second field is assigned by the user. Each PE router advertises its VPN-IPv4 routes with the route target (as one of the BGP path attributes) configured for the VRF table. The route target attached to the advertised route is referred to as the export route target. On the receiving PE router, the route target attached to the route is compared to the route target configured for the local VRF tables. The locally configured route target that is used in deciding whether a VPN-IPv4 route should be installed in a VRF table is referred to as the import route target.

C-Multicast Routing

Customer multicast (C-multicast) routing information exchange refers to the distribution of customer PIM (C-PIM) join/prune messages received from local customer edge (CE) routers to other PE routers (toward the VPN multicast source).

BGP MVPNs

BGP MVPNs use BGP as the control plane protocol between PE routers for MVPNs, including the exchange of C-multicast routing information. The support of BGP as a PE-PE protocol for exchanging C-multicast routes is mandated by Internet draft draft-ietf-l3vpn-mvpn-considerations-06.txt, *Mandatory Features in a Layer 3 Multicast BGP/MPLS VPN Solution*. The use of BGP for distributing C-multicast routing information is closely modeled after its highly successful counterpart of VPN unicast route distribution. Using BGP as the control plane protocol allows service providers to take advantage of this widely deployed, feature-rich protocol. It also enables service providers to leverage their knowledge and investment in managing BGP-MPLS VPN unicast service to offer VPN multicast services.

Sender and Receiver Site Sets

Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt describes an MVPN as a set of administrative policies that determine the PE routers that are in sender and receiver site sets.

A PE router can be a sender, a receiver, or both a sender and a receiver, depending on the configuration:

- A sender site set includes PE routers with local VPN multicast sources (VPN customer multicast sources either directly connected or connected via a CE router). A PE router that is in the sender site set is the sender PE router.
- A receiver site set includes PE routers that have local VPN multicast receivers. A PE router that is in the receiver site set is the receiver PE router.

Provider Tunnels

Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt defines provider tunnels as the transport mechanisms used for forwarding VPN multicast traffic across service provider networks. Different tunneling technologies, such as generic routing encapsulation (GRE) and MPLS, can be used to create provider tunnels. Provider tunnels can be signaled by a variety of signaling protocols. This topic describes only PIM-SM (ASM) signaled IP GRE provider tunnels and RSVP-Traffic Engineering (RSVP-TE) signaled MPLS provider tunnels.

In BGP MVPNs, the sender PE router distributes information about the provider tunnel in a BGP attribute called provider multicast service interface (PMSI). By default, all receiver PE routers join and become the leaves of the provider tunnel rooted at the sender PE router.

Provider tunnels can be inclusive or selective:

- An inclusive provider tunnel (I-PMSI provider tunnel) enables a PE router that is in the sender site set of an MVPN to transmit multicast data to all PE routers that are members of that MVPN.
- A selective provider tunnel (S-PMSI provider tunnel) enables a PE router that is in the sender site set of an MVPN to transmit multicast data to a subset of the PE routers.

RELATED DOCUMENTATION

[Understanding Next-Generation MVPN Network Topology | 716](#)

[Generating Next-Generation MVPN VRF Import and Export Policies Overview | 737](#)

[Exchanging C-Multicast Routes | 1023](#)

Example: Configuring MBGP Multicast VPNs

Understanding Next-Generation MVPN Control Plane

IN THIS SECTION

- [BGP MCAST-VPN Address Family and Route Types | 721](#)
- [Intra-AS MVPN Membership Discovery \(Type 1 Routes\) | 724](#)
- [Inter-AS MVPN Membership Discovery \(Type 2 Routes\) | 725](#)
- [Selective Provider Tunnels \(Type 3 and Type 4 Routes\) | 725](#)
- [Source Active Autodiscovery Routes \(Type 5 Routes\) | 725](#)

- C-Multicast Route Exchange (Type 6 and Type 7 Routes) | 725
- PMSI Attribute | 725
- VRF Route Import and Source AS Extended Communities | 726
- Converting MVPN Type 5 Routes to MSDP SA | 727

The BGP next-generation multicast virtual private network (MVPN) control plane, as specified in Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt and Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, distributes all the necessary information to enable end-to-end C-multicast routing exchange via BGP. The main tasks of the control plane (Table 24 on page 721) include MVPN autodiscovery, distribution of provider tunnel information, and PE-PE C-multicast route exchange.

Table 24: Next-Generation MVPN Control Plane Tasks

Control Plane Task	Description
MVPN autodiscovery	A provider edge (PE) router discovers the identity of the other PE routers that participate in the same MVPN.
Distribution of provider tunnel information	A sender PE router advertises the type and identifier of the provider tunnel that it will use to transmit VPN multicast packets.
PE-PE C-Multicast route exchange	A receiver PE router propagates C-multicast join messages (C-joins) received over its VPN interface toward the VPN multicast sources.

BGP MCAST-VPN Address Family and Route Types

Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt introduced a BGP address family called MCAST-VPN for supporting next-generation MVPN control plane operations. The new address family is assigned the subsequent address family identifier (SAFI) of 5 by the Internet Assigned Numbers Authority (IANA).

A PE router that participates in a BGP-based next-generation MVPN network is required to send a BGP update message that contains MCAST-VPN network layer reachability information (NLRI). An MCAST-VPN NLRI contains route type, length, and variable fields. The value of each variable field depends on the route type.

Seven types of next-generation MVPN BGP routes (also referred to as routes in this topic) are specified (Table 25 on page 722). The first five route types are called autodiscovery MVPN routes. This topic also refers to Type 1-5 routes as non-C-multicast MVPN routes. Type 6 and Type 7 routes are called C-multicast MVPN routes.

Table 25: Next-generation MVPN BGP Route Types

Usage	Type	Name	Description
Membership autodiscovery routes for inclusive provider tunnels	1	Intra autonomous system (intra-AS) I-PMSI autodiscovery route	<ul style="list-style-type: none"> • Originated by all next-generation MVPN PE routers. • Used for advertising and learning intra autonomous system (intra-AS) MVPN membership information.
	2	Inter-AS I-PMSI AD route	<ul style="list-style-type: none"> • Originated by next-generation MVPN ASBR routers. • Used for advertising and learning inter-AS MVPN membership information.
Autodiscovery routes for selective provider tunnels	3	S-PMSI AD route	<ul style="list-style-type: none"> • Originated by a sender router. • Used for initiating a selective provider tunnel for a particular (C-S, C-G).

Table 25: Next-generation MVPN BGP Route Types (Continued)

Usage	Type	Name	Description
	4	Leaf AD route	<ul style="list-style-type: none"> • Originated by receiver PE routers in response to receiving a Type 3 route. • Used by a sender PE router to discover the leaves of a selective provider tunnel. • Also used for inter-AS operations that are not covered in this topic.
VPN multicast source discovery routes	5	Source active AD route	<ul style="list-style-type: none"> • Originated by the PE router that discovers an active VPN multicast source. • Used by PE routers to learn the identity of active VPN multicast sources.
C-Multicast routes	6	Shared tree join route	<ul style="list-style-type: none"> • Originated by receiver PE routers. • Originated when a PE router receives a shared tree C-join (C-*, C-G) through its PE-CE interface.

Table 25: Next-generation MVPN BGP Route Types (Continued)

Usage	Type	Name	Description
	7	Source tree join route	<ul style="list-style-type: none"> • Originated by receiver PE routers. • Originated when a PE router receives a source tree C-join (C-S, C-G) or originated by the PE router that already has a Type 6 route and receives a Type 5 route.

Intra-AS MVPN Membership Discovery (Type 1 Routes)

All next-generation MVPN PE routers create and advertise a Type 1 intra-AS autodiscovery route (Figure 97 on page 724) for each MVPN to which they are connected. Table 26 on page 724 describes the format of each MVPN Type 1 intra-AS autodiscovery route.

Figure 97: Intra-AS I-PMSI AD Route Type MCAST-VPN NLRI Format



8041539

Table 26: Type 1 Intra-AS Autodiscovery Route MVPN Format Descriptions

Field	Description
Route Distinguisher	Set to the route distinguisher configured for the VPN.
Originating Router's IP Address	Set to the IP address of the router originating this route. The address is typically the primary loopback address of the PE router.

Inter-AS MVPN Membership Discovery (Type 2 Routes)

Type 2 routes are used for membership discovery between PE routers that belong to different autonomous systems (ASs). Their use is not covered in this topic.

Selective Provider Tunnels (Type 3 and Type 4 Routes)

A sender PE router that initiates a selective provider tunnel is required to originate a Type 3 intra-AS S-PMSI autodiscovery route with the appropriate PMSI attribute.

A receiver PE router responds to a Type 3 route by originating a Type 4 leaf autodiscovery route if it has local receivers interested in the traffic transmitted on the selective provider tunnel. Type 4 routes inform the sender PE router of the leaf PE routers.

Source Active Autodiscovery Routes (Type 5 Routes)

Type 5 routes carry information about active VPN sources and the groups to which they are transmitting data. These routes can be generated by any PE router that becomes aware of an active source. Type 5 routes apply only for PIM-SM (ASM) when intersite source-tree-only mode is being used.

C-Multicast Route Exchange (Type 6 and Type 7 Routes)

The C-multicast route exchange between PE routers refers to the propagation of C-joins from receiver PE routers to the sender PE routers.

In a next-generation MVPN, C-joins are translated into (or encoded as) BGP C-multicast MVPN routes and advertised via the BGP MCAST-VPN address family toward the sender PE routers.

Two types of C-multicast MVPN routes are specified:

- Type 6 C-multicast routes are used in representing information contained in a shared tree (C-*, C-G) join.
- Type 7 C-multicast routes are used in representing information contained in a source tree (C-S, C-G) join.

PMSI Attribute

The provider multicast service interface (PMSI) attribute ([Figure 98 on page 726](#)) carries information about the provider tunnel. In a next-generation MVPN network, the sender PE router sets up the provider tunnel, and therefore is responsible for originating the PMSI attribute. The PMSI attribute can be attached to Type 1, Type 2, or Type 3 routes. [Table 27 on page 726](#) describes each PMSI attribute format.

Figure 98: PMSI Tunnel Attribute Format

Flags	1 octet
Tunnel Type	1 octet
MPLS Label	3 octets
Tunnel Identifier	Variable

8041540

Table 27: PMSI Tunnel Attribute Format Descriptions

Field	Description
Flags	Currently has only one flag specified: Leaf Information Required. This flag is used for S-PMSI provider tunnel setup.
Tunnel Type	Identifies the tunnel technology used by the sender. Currently there are seven types of tunnels supported.
MPLS Label	Used when the sender PE router allocates the MPLS labels (also called upstream label allocation). This technique is described in RFC 5331 and is outside the scope of this topic.
Tunnel Identifier	Uniquely identifies the tunnel. Its value depends on the value set in the tunnel type field.

For example, Router PE1 originates the following PMSI attribute:

```
PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[10.1.1.1:0:6574:10.1.1.1]
```

VRF Route Import and Source AS Extended Communities

Two extended communities are specified to support next-generation MVPNs: source AS (src-as) and VRF route import extended communities.

The source AS extended community is an AS-specific extended community that identifies the AS from which a route originates. This community is mostly used for inter-AS operations, which is not covered in this topic.

The VPN routing and forwarding (VRF) route import extended community is an IP-address-specific extended community that is used for importing C-multicast routes in the VRF table of the active sender PE router to which the source is attached.

Each PE router creates a unique route target import and src-as community for each VPN and attaches them to the VPN-IPv4 routes.

Converting MVPN Type 5 Routes to MSDP SA

You can convert next-generation multicast virtual private network (MVPN) Type 5 routes to Multicast Source Discovery Protocol (MSDP) source active (SA) routes. This conversion helps in reducing the number of MSDP sessions running between VPN customer rendezvous points (C-RPs). For example, instead of having MSDP running among all C-RPs in a deployment, the C-RPs could instead run their MSDP sessions on a single PE router configured for multiple MSDP peers. The PE router, acting as a C-RP device, receives MVPN SA Type 5 routes from the RP-PE or the source PE router, converts those routes to MSDP, and advertises the MSDP routes to its MSDP peers.

MVPN Type 5 SA routes are added to the MVPN table. These MVPN Type 5 SA routes also include a new Extended Community (EC) with the IPv4 address of the RP where the MVPN SA was generated. The Type 5 route's source and EC are additionally added to the MSDP table. Stale routes, including the EC, are removed from the MSDP table after the MVPN Type 5 SA route is cleared from the MVPN table.

Conversion of MVPN Type 5 routes to MSDP SA is defined in [RFC 9081](#). Previously, Junos OS only supported conversion from MSDP SA routes to MVPN Type 5.



NOTE: MVPN Type 5 routes to MSDP SA conversion works in spt-only mode. rpt-spt mode is not supported.

The Extended Community is IPv4 only. IPv6 is not supported, as MSDP is IPv4 only.

Benefits of converting MVPN Type 5 routes to MSDP SA

- For MVPN provider networks, conversion of MVPN Type 5 routes to MSDP SA eliminates the need for VPN-specific MSDP sessions required among the Provider Edge (PE) routers that are customer MSDP peers.

Enabling the conversion of MVPN Type 5 routes to MSDP SA

To enable the conversion of Type 5 routes to MSDP SA, use the `set convert-sa-to-msdp` configuration statement at the hierarchy level shown below:

```
[edit routing-instance instance-name protocols mvpn mvpn-mode spt-only]
user@router# set convert-sa-to-msdp
```

Verifying the conversion of MVPN Type 5 routes to MSDP SA

Purpose

Verify that the MSDP table contains source active MVPN Type 5 routes.

Action

From operational mode, run the following show commands:

- `show route table instance-name.mvpn.0 extensive`—View and establish the MVPN Type 5 route.
- `show route table instance-name.inet.4`—View the SA route converted from the MVPN Type 5 route.
- `show msdp source-active`—Verify that the MVPN Type 5 route is populated in the MSDP table.
- `show msdp source-active instance instance-name`—Verify that the MVPN Type 5 route is populated in the specific instance of the MSDP table.

```
user@host> show route table VPN-A.mvpn.0 extensive
5:192.168.1.1:1:32:10.0.0.2:32:233.252.0.1/240 (1 entry, 1 announced)
 *PIM Preference: 105
Next hop type: Multicast (IPv4) Composite, Next hop index: 0
Address: 0xa3d7c24
Next-hop reference count: 3
State: <Active Int>
Age: 18
Validation State: unverified
Task: PIM.VPN-A
Announcement bits (3): 0-PIM.VPN-A 1-mvpn global task 2-rt-export
```

```
AS path: I
Communities: mvpn-sa-rp:192.168.0.1:0
```

```
user@host> show route table instance.inet.4
instance.inet.4: 1 destination, 1 route (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

233.252.0.1,10.0.0.2/64*[MSDP/175/2] 00:00:46, from 192.168.0.1
Local
```

```
user@host> show msdp source-active
Group address Source address Peer address Originator Flags
233.252.0.1 10.0.0.2 192.168.0.3 192.168.0.1 Accept
```

```
user@host> show msdp source-active instance VPN-A
Group address Source address Peer address Originator Flags
233.252.0.1 10.0.0.2 local 192.168.0.1 Accept
```

Meaning

- The `show route table VPN-A.mvpn.0` extensive output shows the Extended community (EC) with the IPv4 address of the RP where the MVPN SA was generated, under the Type 5 route table.
- The `show route table instance.inet.4` output shows the route table with the multicast group address, the source and the SA converted route.
- The `show msdp source-active` output shows the MSDP table with all the source active routes received from its neighbors. In this case, 192.168.0.1 is displayed as the originator, which is the IPv4 address of the source RP received from the Type 5 route.
- The `show msdp source-active instance VPN-A` output shows the MSDP table for VPN-A. In this case too, 192.168.0.1 is displayed as the originator, which is the IPv4 address of the source RP received from the Type 5 route.

RELATED DOCUMENTATION

[Next-Generation MVPN Data Plane Overview](#) | 730

[Distributing C-Multicast Routes Overview | 1017](#)

[Enabling Next-Generation MVPN Services | 734](#)

[Signaling Provider Tunnels and Data Plane Setup | 1037](#)

[Originating Type 1 Intra-AS Autodiscovery Routes Overview | 1033](#)

[Understanding Next-Generation MVPN Network Topology | 716](#)

Next-Generation MVPN Data Plane Overview

IN THIS SECTION

- [Inclusive Provider Tunnels | 731](#)
- [Selective Provider Tunnels \(S-PMSI Autodiscovery/Type 3 and Leaf Autodiscovery/Type 4 Routes\) | 732](#)

A next-generation multicast virtual private network (MVPN) data plane is composed of provider tunnels originated by and rooted at the sender provider edge (PE) routers and the receiver PE routers as the leaves of the provider tunnel.

A provider tunnel can carry data for one or more VPNs. Those provider tunnels that carry data for more than one VPN are called aggregate provider tunnels and are outside the scope of this topic. Here, we assume that a provider tunnel carries data for only one VPN.

This topic covers two types of tunnel technologies: IP generic routing encapsulation (GRE) provider tunnels signaled by Protocol Independent Multicast-Sparse Mode (PIM-SM) any-source multicast (ASM) and MPLS provider tunnels signaled by RSVP-Traffic Engineering (RSVP-TE).

When a provider tunnel is signaled by PIM, the sender PE router runs another instance of the PIM protocol on the provider's network (P-PIM) that signals a provider tunnel for that VPN. When a provider tunnel is signaled by RSVP-TE, the sender PE router initiates a point-to-multipoint label-switched path (LSP) toward receiver PE routers by using point-to-multipoint RSVP-TE protocol messages. In either case, the sender PE router advertises the tunnel signaling protocol and the tunnel ID to other PE routers via BGP by attaching the provider multicast service interface (PMSI) attribute to either the Type 1 intra-AS autodiscovery routes (inclusive provider tunnels) or Type 3 S-PMSI autodiscovery routes (selective provider tunnels).



NOTE: The sender PE router goes through two steps when setting up the data plane. First, using the PMSI attribute, it advertises the provider tunnel it is using via BGP. Second, it actually signals the tunnel using whatever tunnel signaling protocol is configured for that VPN. This allows receiver PE routers to bind the tunnel that is being signaled to the VPN that imported the Type 1 intra-AS autodiscovery route. Binding a provider tunnel to a VRF table enables a receiver PE router to map the incoming traffic from the core network on the provider tunnel to the local target VRF table.

The PMSI attribute contains the provider tunnel type and an identifier. The value of the provider tunnel identifier depends on the tunnel type. [Table 28 on page 731](#) identifies the tunnel types specified in Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt.

Table 28: Tunnel Types Supported by PMSI Tunnel Attribute

Tunnel Type	Description
0	No tunnel information present
1	RSVP-TE point-to-multipoint LSP
2	Multicast LDP point-to-multipoint LSP
3	PIM-SSM tree
4	PIM-SM tree
5	PIM-Bidir tree
6	Ingress replication
7	Multicast LDP multipoint-to-multipoint LSP

Inclusive Provider Tunnels

This section describes various types of provider tunnels and attributes of provider tunnels.

PMSI Attribute of Inclusive Provider Tunnels Signaled by PIM-SM

When the Tunnel Type field of the PMSI attribute is set to 4 (PIM-SM Tree), the tunnel identifier field contains <Sender Address, P-Multicast Group Address>. The Sender Address field is set to the router ID of the sender PE router. The P-multicast group address is set to a multicast group address from the service provider's P-multicast address space and uniquely identifies the VPN. A receiver PE router that receives an intra-AS autodiscovery route with a PMSI attribute whose tunnel type is PIM-SM is required to join the provider tunnel.

For example, if the service provider deploys PIM-SM provider tunnels (instead of RSVP-TE provider tunnels), Router PE1 advertises the following PMSI attribute:

```
PMSI: 0:PIM-SM:label[0:0:0]:Sender10.1.1.1 Group 239.1.1.1
```

PMSI Attribute of Inclusive Provider Tunnels Signaled by RSVP-TE

When the tunnel type field of the PMSI attribute is set to 1 (RSVP-TE point-to-multipoint LSP), the tunnel identifier field contains an RSVP-TE point-to-multipoint session object as described in RFC 4875. The session object contains the <Extended Tunnel ID, Reserved, Tunnel ID, P2MP ID> associated with the point-to-multipoint LSPs.

The PE router that originates the PMSI attribute is required to signal an RSVP-TE point-to-multipoint LSP and the sub-LSPs. A PE router that receives this PMSI attribute must establish the appropriate state to properly handle the traffic received over the sub-LSP.

For example, Router PE1 advertises the following PMSI attribute:

```
PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[10.1.1.1:0:6574:10.1.1.1]
```

Selective Provider Tunnels (S-PMSI Autodiscovery/Type 3 and Leaf Autodiscovery/Type 4 Routes)

A selective provider tunnel is used for mapping a specific C-multicast flow (a (C-S, C-G) pair) onto a specific provider tunnel. There are a variety of situations in which selective provider tunnels can be useful. For example, they can be used for putting high-bandwidth VPN multicast data traffic onto a separate provider tunnel rather than the default inclusive provider tunnel, thus restricting the distribution of traffic to only those PE routers with active receivers.

In BGP next-generation multicast virtual private networks (MVPNs), selective provider tunnels are signaled using Type 3 Selective-PMSI (S-PMSI) autodiscovery routes. See [Figure 99 on page 733](#) and [Table 29 on page 733](#) for details. The sender PE router sends a Type 3 route to signal that it is sending traffic for a particular (C-S, C-G) flow using an S-PMSI provider tunnel.

Figure 99: S-PMSI Autodiscovery Route Type Multicast (MCAST)-VPN Network Layer Reachability Information (NLRI) Format

Route Distinguisher	8 octets
Multicast Source Length	1 octet
Multicast Source	Variable
Multicast Group Length	1 octet
Multicast Group	Variable

8041544

Table 29: S-PMSI Autodiscovery Route Type Format Descriptions

Field	Description
Route Distinguisher	Set to the route distinguisher configured on the router originating this route.
Multicast Source Length	Set to 32 for IPv4 and to 128 for IPv6 C-S IP addresses.
Multicast Source	Set to the C-S IP address.
Multicast Group Length	Set to 32 for IPv4 and to 128 for IPv6 C-G addresses.
Multicast Group	Set to the C-G address.

The S-PMSI autodiscovery (Type 3) route carries a PMSI attribute similar to the PMSI attribute carried with intra-AS autodiscovery (Type 1) routes. The `Flags` field of the PMSI attribute carried by the S-PMSI autodiscovery route is set to the leaf information required. This flag signals receiver PE routers to originate a Type 4 leaf autodiscovery route ([Figure 100 on page 734](#)) to join the selective provider tunnel if they have active receivers. See [Table 30 on page 734](#) for details of leaf autodiscovery route type MCAST-VPN NLRI format descriptions.

Figure 100: Leaf Autodiscovery Route Type MCAST-VPN NLRI Format



8041545

Table 30: Leaf Autodiscovery Route Type MCAST-VPN NLRI Format Descriptions

Field	Description
Route Key	Contains the original Type 3 route received.
Originating Router's IP Address	Set to the IP address of the PE router originating the leaf autodiscovery route. This is typically the primary loopback address.

RELATED DOCUMENTATION

[Understanding Next-Generation MVPN Control Plane | 720](#)

[Enabling Next-Generation MVPN Services | 734](#)

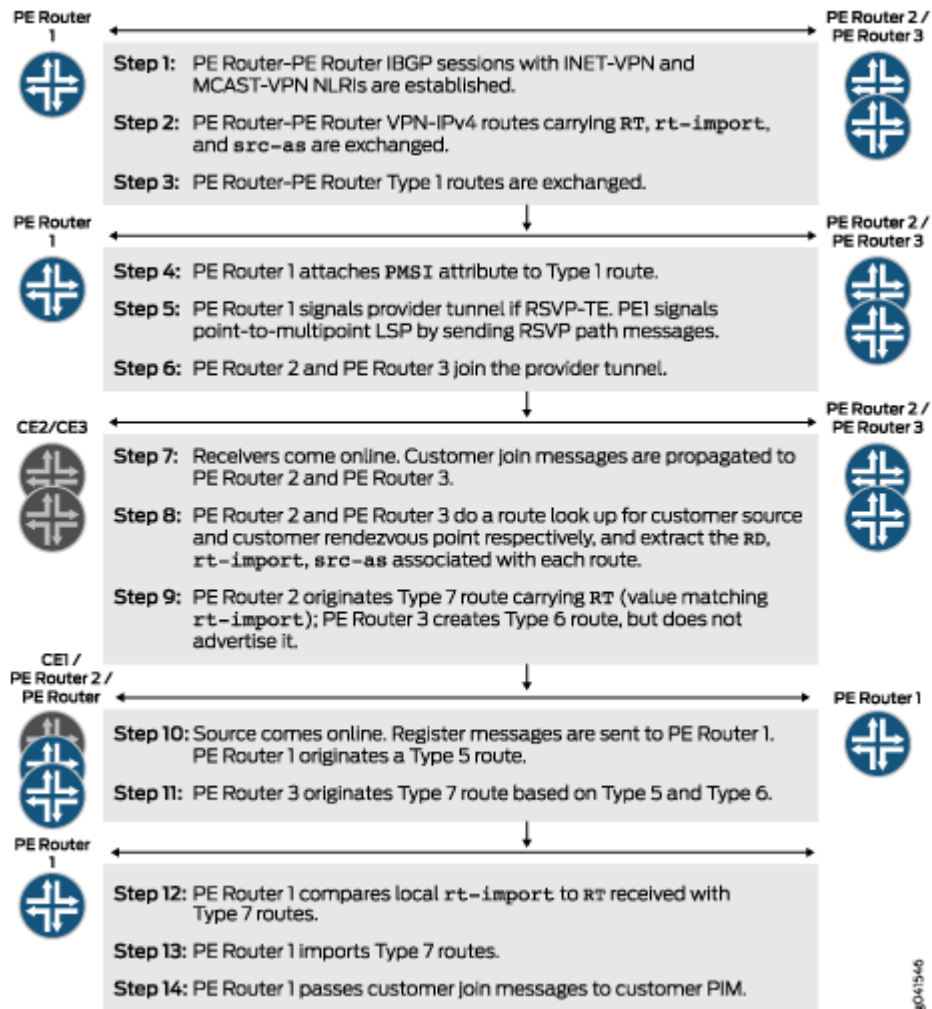
[Signaling Provider Tunnels and Data Plane Setup | 1037](#)

[Understanding Next-Generation MVPN Network Topology | 716](#)

Enabling Next-Generation MVPN Services

Juniper Networks introduced the industry's first implementation of BGP next-generation multicast virtual private networks (MVPNs). See [Figure 101 on page 735](#) for a summary of a Junos OS next-generation MVPN routing flow.

Figure 101: Junos OS Next-Generation MVPN Routing Flow



Next-generation MVPN services are configured on top of BGP-MPLS unicast VPN services.

You can configure a Juniper Networks PE router that is already providing unicast BGP-MPLS VPN connectivity to support multicast VPN connectivity in three steps:

1. Configure the provider edge (PE) routers to support the BGP multicast VPN address family by including the signaling statement at the `[edit protocols bgp group group-name family inet-mvpn]` hierarchy level. This address family enables PE routers to exchange MVPN routes.
2. Configure the PE routers to support the MVPN control plane tasks by including the `mvpn` statement at the `[edit routing-instances routing-instance-name protocols]` hierarchy level. This statement signals PE routers to initialize the MVPN module that is responsible for the majority of next-generation MVPN control plane tasks.
3. Configure the sender PE router to signal a provider tunnel by including the `provider-tunnel` statement at the `[edit routing-instances routing-instance-name]` hierarchy level. You must also enable the tunnel

signaling protocol (RSVP-TE or P-PIM) if it is not part of the unicast VPN service configuration. To enable the tunnel signaling protocol, include the `rsvp-te` or `pim-asm` statements at the [edit routing-instances *routing-instance-name* provider-tunnel] hierarchy level.

After these three statements are configured and each PE router has established internal BGP (IBGP) sessions using both INET-VPN and MCAST-VPN address families, four routing tables are automatically created. These tables are `bgp.l3vpn.0`, `bgp.mvpn.0`, `<routing-instance-name>.inet.0`, and `<routing-instance-name>.mvpn.0`. See [Table 31 on page 736](#)

Table 31: Automatically Generated Routing Tables

Automatically Generated Routing Table	Description
<code>bgp.l3vpn.0</code>	Populated with VPN-IPv4 routes received from remote PE routers via the INET-VPN address family. The routes in the <code>bgp.l3vpn.0</code> table are in the form of RD:IPv4-address and carry one or more routing table communities. In a next-generation MVPN network, these routes also carry <code>rt-import</code> and <code>src-as</code> communities.
<code>bgp.mvpn.0</code>	Populated by MVPN routes (Type 1 – Type 7). Received from remote PE routers via the MCAST-VPN address family. Routes in this table carry one or more routing table communities.
<code><routing-instance-name>.inet.0</code>	Populated by local and remote VPN unicast routes. The local VPN routes are typically learned from local CE routers via protocols such as BGP, OSPF, and RIP, or via a static configuration. The remote VPN routes are imported from the <code>bgp.l3vpn.0</code> table if their routing table matches one of the import routing tables configured for the VPN. When remote VPN routes are imported from the <code>bgp.l3vpn.0</code> table, their route distinguisher is removed, leaving them as regular unicast IPv4 addresses.

Table 31: Automatically Generated Routing Tables (Continued)

Automatically Generated Routing Table	Description
<routing-instance-name>.mvpn.0	Populated by local and remote MVPN routes. The local MVPN routes are typically the locally originated routes, such as Type 1 intra-AS autodiscovery routes, or Type 7 C-multicast routes. The remote MVPN routes are imported from the <code>bgp.mvpn.0</code> table based on their route target. The import route target used for accepting MVPN routes into the <routing-instance-name>.mvpn.0 table is different for C-multicast MVPN routes (Type 6 and Type 7) versus non-C-multicast MVPN routes (Type 1 – Type 5).

RELATED DOCUMENTATION

[Understanding Next-Generation MVPN Network Topology | 716](#)

[Generating Next-Generation MVPN VRF Import and Export Policies Overview | 737](#)

[Generating Source AS and Route Target Import Communities Overview | 1032](#)

[Originating Type 1 Intra-AS Autodiscovery Routes Overview | 1033](#)

[Signaling Provider Tunnels and Data Plane Setup | 1037](#)

Generating Next-Generation MVPN VRF Import and Export Policies Overview

IN THIS SECTION

- [Policies That Support Unicast BGP-MPLS VPN Services | 738](#)
- [Policies That Support Next-Generation MVPN Services | 739](#)

In Junos OS, the policy module is responsible for VPN routing and forwarding (VRF) route import and export decisions. You can configure these policies explicitly, or Junos OS can generate them internally for you to reduce user-configured statements and simplify configuration. Junos OS generates all

necessary policies for supporting next-generation multicast virtual private network (MVPN) import and export decisions. Some of these policies affect normal VPN unicast routes.

The system gives a name to each internal policy it creates. The name of an internal policy starts and ends with a “__” notation. Also the keyword `internal` is added at the end of each internal policy name. You can display these internal policies using the `show policy` command.

Policies That Support Unicast BGP-MPLS VPN Services

A Juniper Networks provider edge (PE) router requires a `vrf-import` and a `vrf-export` policy to control unicast VPN route import and export decisions for a VRF. You can configure these policies explicitly at the `[edit routing-instances routing-instance-name vrf-import import_policy_name]` and `[edit routing-instances routing-instance-name vrf-export export_policy_name]` hierarchy level. Alternately, you can configure only the route target for the VRF at the `[edit routing-instances routing-instance-name vrf-target]` hierarchy level, and Junos OS then generates these policies automatically for you. Routers referenced in this topic are shown in ["Understanding Next-Generation MVPN Network Topology" on page 716](#).

The following list identifies the automatically generated policy names and where they are applied:

Policy: `vrf-import`

Naming convention: `__vrf-import-<routing-instance-name>-internal__`

Applied to: VPN-IPv4 routes in the `bgp.l3vpn.0` table

Policy: `vrf-export`

Naming convention: `__vrf-export-<routing-instance-name>-internal__`

Applied to: Local VPN routes in the `<routing-instance-name>.inet.0` table

Use the `show policy __vrf-import-vpna-internal__` command to verify that Router PE1 has created the following `vrf-import` and `vrf-export` policies based on a `vrf-target` of `target:10:1`. In this example, we see that the `vrf-import` policy is constructed to accept a route if the route target of the route matches `target:10:1`. Similarly, a route is exported with a route target of `target:10:1`.

```
user@PE1> show policy __vrf-import-vpna-internal__
Policy __vrf-import-vpna-internal__:
  Term unnamed:
    from community __vrf-community-vpna-common-internal__ [target:10:1]
    then accept
  Term unnamed:
    then reject
user@PE1> show policy __vrf-export-vpna-internal__
Policy __vrf-export-vpna-internal__:
```

```
Term unnamed:
      then community + __vrf-community-vpna-common-internal__ [target:10:1] accept
```

The values in this example are as follows:

- Internal import policy name: __vrf-import-vpna-internal__
- Internal export policy name: __vrf-export-vpna-internal__
- RT community used in both import and export policies: __vrf-community-vpna-common-internal__
- RT value: target:10:1

Policies That Support Next-Generation MVPN Services

When you configure the `mvpn` statement at the [edit `routing-instances <routing-instance-name> protocols`] hierarchy level, Junos OS automatically creates three new internal policies: one for export, one for import, and one for handling Type 4 routes. Routers referenced in this topic are shown in ["Understanding Next-Generation MVPN Network Topology" on page 716](#).

The following list identifies the automatically generated policy names and where they are applied:

Policy 1: This policy is used to attach `rt-import` and `src-as` extended communities to VPN-IPv4 routes.

Policy name: __vrf-mvpn-export-inet-<routing-instance-name>-internal__

Applied to: All routes in the <routing-instance-name>inet.0 table

Use the `show policy __vrf-mvpn-export-inet-vpna-internal__` command to verify that the following export policy is created on Router PE1. Router PE1 adds `rt-import:10.1.1.1:64` and `src-as:65000:0` communities to unicast VPN routes through this policy.

```
user@PE1> show policy __vrf-mvpn-export-inet-vpna-internal__
Policy __vrf-mvpn-export-inet-vpna-internal__:
  Term unnamed:
    then community + __vrf-mvpn-community-rt_import-vpna-internal__ [rt-
import:10.1.1.1:64 ] community + __vrf-mvpn-community-src_as-vpna-internal__ [src-as:65000:0 ]
accept
```

The values in this example are as follows:

- Policy name: __vrf-mvpn-export-inet-vpna-internal__
- `rt-import` community name: __vrf-mvpn-community-rt_import-vpna-internal__
- `rt-import` community value: rt-import:10.1.1.1:64

- src-as community name: `__vrf-mvpn-community-src_as-vpna-internal__`
- src-as community value: `src-as:65000:0`

Policy 2: This policy is used to import C-Mmulticast routes from the `bgp.mvpn.0` table to the `<routing-instance-name>.mvpn.0` table.

Policy name: `__vrf-mvpn-import-cmcast-<routing-instance-name>-internal__`

Applied to: C-multicast (MVPN) routes in the `bgp.mvpn.0` table

Use the `show policy __vrf-mvpn-import-cmcast-vpna-internal__` command to verify that the following import policy is created on Router PE1. The policy accepts those C-multicast MVPN routes carrying a route target of `target:10.1.1.1:64` and installs them in the `vpna.mvpn.0` table.

```
user@PE1> show policy __vrf-mvpn-import-cmcast-vpna-internal__
Policy __vrf-mvpn-import-cmcast-vpna-internal__:
  Term unnamed:
    from community __vrf-mvpn-community-rt_import-target-vpna-internal__
  [target:10.1.1.1:64 ]
    then accept
  Term unnamed:
    then reject
```

The values in this example are as follows:

- Policy name: `__vrf-mvpn-import-cmcast-vpna-internal__`
- C-multicast import RT community: `__vrf-mvpn-community-rt_import-target-vpna-internal__`
- Community value: `target:10.1.1.1:64`

Policy 3: This policy is used for importing Type 4 routes and is created by default even if a selective provider tunnel is not configured. The policy affects only Type 4 routes received from receiver PE routers.

Policy name: `__vrf-mvpn-import-cmcast-leafAD-global-internal__`

Applied to: Type 4 routes in the `bgp.mvpn.0` table

Use the `show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__` command to verify that the following import policy is created on Router PE1.

```
user@PE1> show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__
Policy __vrf-mvpn-import-cmcast-leafAD-global-internal__:
  Term unnamed:
```

```

        from community __vrf-mvpn-community-rt_import-target-global-internal__
    [target:10.1.1.1:0 ]
        then accept
    Term unnamed:
        then reject

```

RELATED DOCUMENTATION

[Understanding MBGP Multicast VPN Extranets | 858](#)

[Example: Configuring MBGP Multicast VPN Extranets | 860](#)

Example: Configuring MBGP Multicast VPNs

[Enabling Next-Generation MVPN Services | 734](#)

Multiprotocol BGP MVPNs Overview

IN THIS SECTION

- [Comparison of Draft Rosen Multicast VPNs and Next-Generation Multiprotocol BGP Multicast VPNs | 741](#)
- [MBGP Multicast VPN Sites | 742](#)
- [Multicast VPN Standards | 743](#)
- [PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs | 743](#)
- [MBGP-Based Multicast VPN Trees | 744](#)

Comparison of Draft Rosen Multicast VPNs and Next-Generation Multiprotocol BGP Multicast VPNs

There are several multicast applications driving the deployment of next-generation Layer 3 multicast VPNs (MVPNs). Some of the key emerging applications include the following:

- Layer 3 VPN multicast service offered by service providers to enterprise customers
- Video transport applications for wholesale IPTV and multiple content providers attached to the same network
- Distribution of media-rich financial services or enterprise multicast services

- Multicast backhaul over a metro network

There are two ways to implement Layer 3 MVPNs. They are often referred to as dual PIM MVPNs (also known as “draft-rosen”) and multiprotocol BGP (MBGP)-based MVPNs (the “next generation” method of MVPN configuration). Both methods are supported and equally effective. The main difference is that the MBGP-based MVPN method does not require multicast configuration on the service provider backbone. Multiprotocol BGP multicast VPNs employ the intra-autonomous system (AS) next-generation BGP control plane and PIM sparse mode as the data plane. The PIM state information is maintained between the PE routers using the same architecture that is used for unicast VPNs. The main advantage of deploying MVPNs with MBGP is simplicity of configuration and operation because multicast is not needed on the service provider VPN backbone connecting the PE routers.

Using the draft-rosen approach, service providers might experience control and data plane scaling issues associated with the maintenance of two routing and forwarding mechanisms: one for VPN unicast and one for VPN multicast. For more information on the limitations of Draft Rosen, see draft-rekhter-mboned-mvpn-deploy.

SEE ALSO

| No Link Title

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.

SEE ALSO

[Example: Allowing MBGP MVPN Remote Sources](#)

[Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN](#)

Multicast VPN Standards

MBGP MVPNs are defined in the following IETF Internet drafts:

- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-03.txt, *BGP Encodings for Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-02.txt, *Multicast in MPLS/BGP IP VPNs*

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.

SEE ALSO

[Example: Allowing MBGP MVPN Remote Sources](#)

[Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN](#)

MBGP-Based Multicast VPN Trees

MBGP-based MVPNs (next-generation MVPNs) are based on Internet drafts and extend unicast VPNs based on RFC 2547 to include support for IP multicast traffic. These MVPNs follow the same architectural model as the unicast VPNs and use BGP as the provider edge (PE)-to-PE control plane to exchange information. The next generation MVPN approach is based on Internet drafts draft-ietf-l3vpn-2547bis-mcast.txt, draft-ietf-l3vpn-2547bis-mcast-bgp.txt, and draft-morin-l3vpn-mvpn-considerations.txt.

MBGP-based MVPNs introduce two new types of tree:

- | | |
|-----------------------|--|
| Inclusive tree | A single multicast distribution tree in the backbone carrying all the multicast traffic from a specified set of one or more MVPNs. An inclusive tree carrying the traffic of more than one MVPN is an <i>aggregate inclusive tree</i> . All the PEs that attach to MVPN receiver sites using the tree belong to that inclusive tree. |
| Selective tree | A single multicast distribution tree in the backbone carrying traffic for a specified set of one or more multicast groups. When multicast groups belonging to more than one MVPN are on the tree, it is called an <i>aggregate selective tree</i> . |

By default, traffic from most multicast groups can be carried by an inclusive tree, while traffic from some groups (for example, high bandwidth groups) can be carried by one of the selective trees. Selective trees, if they contain only those PEs that need to receive multicast data from one or more groups assigned to the tree, can provide more optimal routing than inclusive trees alone, although this requires more state information in the P routers.

An MPLS-based VPN running BGP with autodiscovery is used as the basis for a next-generation MVPN. The autodiscovered route information is carried in MBGP network layer reachability information (NLRI) updates for multicast VPNs (MCAST-VPNs). These MCAST-VPN NLRIs are handled in the same way as IPv4 routes: route distinguishers are used to distinguish between different VPNs in the network. These NLRIs are imported and exported based on the route target extended communities, just as IPv4 unicast routes. In other words, existing BGP mechanisms are used to distribute multicast information on the provider backbone without requiring multicast directly.

For example, consider a customer running Protocol-Independent Multicast (PIM) sparse mode in source-specific multicast (SSM) mode. Only source tree join customer multicast (c-multicast) routes are required. (PIM sparse mode in anysource multicast (ASM) mode can be supported with a few enhancements to SSM mode.)

The customer multicast route carrying a particular multicast source *S* needs to be imported only into the VPN routing and forwarding (VRF) table on the PE router connected to the site that contains the source *S* and not into any other VRF, even for the same MVPN. To do this, each VRF on a particular PE has a distinct VRF route import extended community associated with it. This community consists of the PE router's IP address and local PE number. Different MVPNs on a particular PE have different route imports, and for a particular MVPN, the VRF instances on different PE routers have different route imports. This VRF route import is auto-configured and not controlled by the user.

Also, all the VRFs within a particular MVPN will have information about VRF route imports for each VRF. This is accomplished by “piggybacking” the VRF route import extended community onto the unicast VPN IPv4 routes. To make sure a customer multicast route carrying multicast source *S* is imported only into the VRF on the PE router connected to the site contained the source *S*, it is necessary to find the unicast VPN IPv4 route to *S* and set the route target of the customer multicast route to the VRF import route carried by the VPN IPv4 route just found.

The process of originating customer multicast routes in an MBGP-based MVPN is shown in [Figure 102 on page 746](#).


In the figure, an MVPN has three receiver sites (R1, R2, and R3) and one source site (S). The site routers are connected to four PE routers, and PIM is running between the PE routers and the site routers. However, only BGP runs between the PE routers on the provider's network.

When router PE-1 receives a PIM join message for (S,G) from site router R1, this means that site R1 has one or more receivers for a given source and multicast group (S,G) combination. In that case, router PE-1 constructs and originates a customer multicast route after doing three things:

1. Finding the unicast VPN IPv4 router to source *S*
2. Extracting the route distinguisher and VRF route import from this route
3. Putting the (S,G) information from the PIM join, the router distinguisher from the VPN IPv4 route, and the route target from the VRF route import of the VPN IPv4 route into a MBGP update

The update is distributed around the VPN through normal BGP mechanisms such as router reflectors.

Figure 102: Source and Receiver Sites in an MVPN

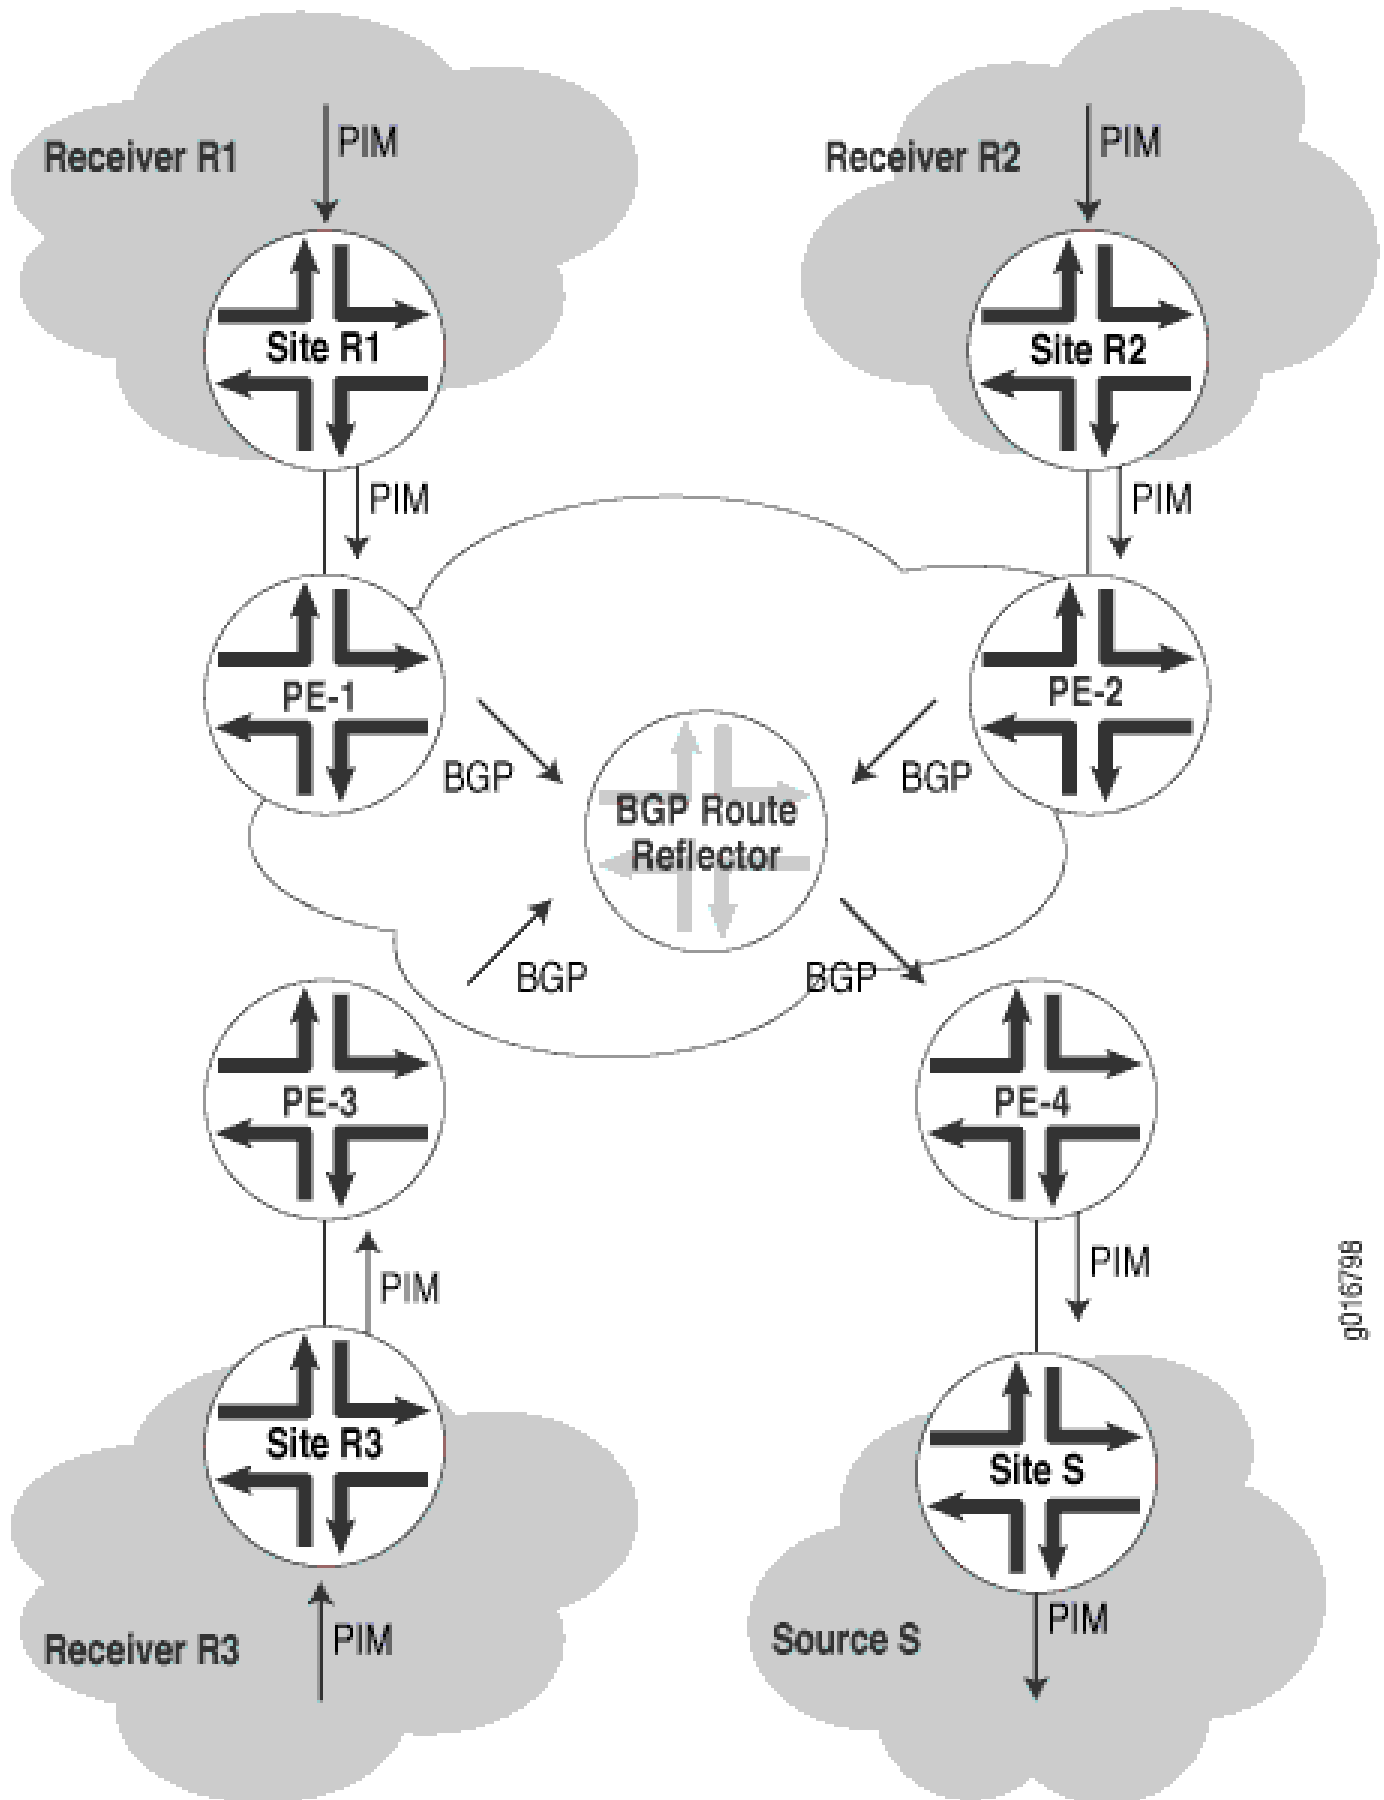


What happens when the source site S receives the MBGP information is shown in [Figure 103 on page 748](#). In the figure, the customer multicast route information is distributed by the BGP route reflector as an MBGP update.

The provider router PE-4 will then:

1. Receive the customer multicast route originated by the PE routers and aggregated by the route reflector.
2. Accept the customer multicast route into the VRF for the correct MVPN (because the VRF route import matches the route target carried in the customer multicast route information).
3. Create the proper (S,G) state in the VRF and propagate the information to the customer routers of source site S using PIM.

Figure 103: Adding a Receiver to an MVPN Source Site Using MBGP



SEE ALSO

| [Example: Configuring Any-Source Multicast for Draft-Rosen VPNs](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

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

| [Configuring Multiprotocol BGP Multicast VPNs | 749](#)

Configuring Multiprotocol BGP Multicast VPNs

IN THIS SECTION

- [Understanding Multiprotocol BGP-Based Multicast VPNs: Next-Generation | 750](#)
- [Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs | 751](#)
- [Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs | 758](#)
- [Example: Configuring MBGP Multicast VPNs | 776](#)
- [Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN | 800](#)
- [Example: Allowing MBGP MVPN Remote Sources | 812](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family | 819](#)
- [Example: Configuring MBGP Multicast VPN Topology Variations | 834](#)
- [Configuring Nonstop Active Routing for BGP Multicast VPN | 850](#)

Understanding Multiprotocol BGP-Based Multicast VPNs: Next-Generation

IN THIS SECTION

- [Route Reflector Behavior in MVPNs | 750](#)

Multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast) constitute the next evolution after dual multicast VPNs (draft-rosen) and provide a simpler solution for administrators who want to configure multicast over Layer 3 VPNs.

The main characteristics of multiprotocol BGP-based multicast VPNs are:

- They extend Layer 3 VPN service (RFC 2547) to support IP multicast for Layer 3 VPN service providers.
- They follow the same architecture as specified by RFC 2547 for unicast VPNs. Specifically, BGP is used as the control plane.
- They eliminate the requirement for the virtual router (VR) model, which is specified in Internet draft draft-rosen-vpn-mcast, *Multicast in MPLS/BGP VPNs*, for multicast VPNs.
- They rely on RFC-based unicast with extensions for intra-AS and inter-AS communication.

Multiprotocol BGP-based VPNs are defined by two sets of sites: a sender set and a receiver set. Hosts within a receiver site set can receive multicast traffic and hosts within a sender site set can send multicast traffic. A site set can be both receiver and sender, which means that hosts within such a site can both send and receive multicast traffic. Multiprotocol BGP-based VPNs can span organizations (so the sites can be intranets or extranets), can span service providers, and can overlap.

Site administrators configure multiprotocol BGP-based VPNs based on customer requirements and the existing BGP and MPLS VPN infrastructure.

Route Reflector Behavior in MVPNs

BGP-based multicast VPN (MVPN) customer multicast routes are aggregated by route reflectors. A route reflector (RR) might receive a customer multicast route with the same NLRI from more than one provider edge (PE) router, but the RR readvertises only one such NLRI. If the set of PE routers that advertise this NLRI changes, the RR does not update the route. This minimizes route churn. To achieve this, the RR sets the next hop to self. In addition, the RR sets the originator ID to itself. The RR avoids unnecessary best-path computation if it receives a subsequent customer multicast route for an NLRI that the RR is already advertising. This allows aggregation of source active and customer multicast routes with the same MVPN NLRI.

SEE ALSO

[Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs | 751](#)

Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

IN THIS SECTION

- [Requirements | 751](#)
- [Overview | 753](#)
- [Configuration | 755](#)
- [Verification | 757](#)

This example shows how to configure point-to-multipoint (P2MP) LDP label-switched paths (LSPs) as the data plane for intra-autonomous system (AS) multiprotocol BGP (MBGP) multicast VPNs (MVPNs). This feature is well suited for service providers who are already running LDP in the MPLS backbone and need MBGP MVPN functionality.

Requirements

Before you begin:

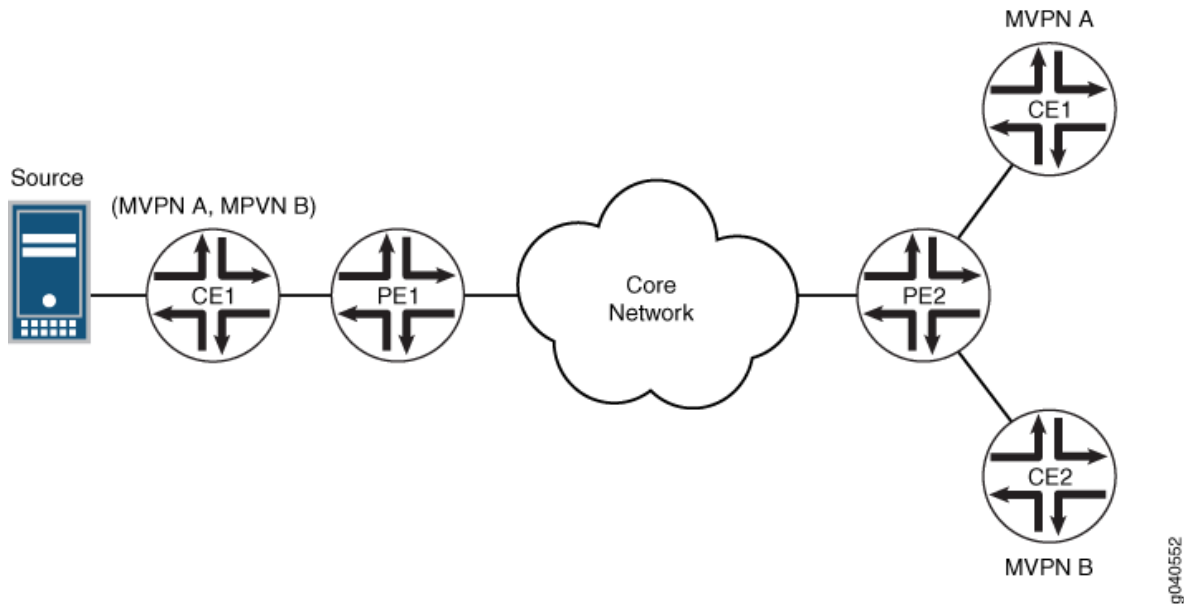
- Configure the router interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure a BGP-MVPN control plane. See "MBGP-Based Multicast VPN Trees" on page 744 in the [Multicast Protocols User Guide](#).
- Configure LDP as the signaling protocol on all P2MP provider and provider-edge routers. See [LDP Operation](#) in the [Junos OS MPLS Applications User Guide](#).
- Configure P2MP LDP LSPs as the provider tunnel technology on each PE router in the MVPN that belongs to the sender site set. See the [Junos OS MPLS Applications User Guide](#).
- Configure either a virtual loopback tunnel interface (requires a Tunnel PIC) or the `vrf-table-label` statement in the MVPN routing instance. If you configure the `vrf-table-label` statement, you can configure an optional virtual loopback tunnel interface as well.

- In an extranet scenario when the egress PE router belongs to multiple MVPN instances, all of which need to receive a specific multicast stream, a virtual loopback tunnel interface (and a Tunnel PIC) is required on the egress PE router. See [Configuring Virtual Loopback Tunnels for VRF Table Lookup](#) in the [Junos OS Services Interfaces Library for Routing Devices](#).
- If the egress PE router is also a transit router for the point-to-multipoint LSP, a virtual loopback tunnel interface (and a Tunnel PIC) is required on the egress PE router. See [Configuring Virtual Loopback Tunnels for VRF Table Lookup](#) in the [Multicast Protocols User Guide](#).
- Some extranet configurations of MBGP MVPNs with point-to-multicast LDP LSPs as the data plane require a virtual loopback tunnel interface (and a Tunnel PIC) on egress PE routers. When an egress PE router belongs to multiple MVPN instances, all of which need to receive a specific multicast stream, the `vrf-table-table` statement cannot be used. In Figure 1, the CE1 and CE2 routers belong to different MVPNs. However, they want to receive a multicast stream being sent by Source. If the `vrf-table-label` statement is configured on Router PE2, the packet cannot be forwarded to both CE1 and CE2. This causes packet loss. The packet is forwarded to both Routers CE1 and CE2 if a virtual loopback tunnel interface is used in both MVPN routing instances on Router PE2. Thus, you need to set up a virtual loopback tunnel interface if you are using an extranet scenario wherein the egress PE router belongs to multiple MVPN instances that receive a specific multicast stream, or if you are using the egress PE router as a transit router for the point-to-multipoint LSP.



NOTE: Starting in Junos OS Release 15.1X49-D50 and Junos OS Release 17.3R1, the `vrf-table-label` statement allows mapping of the inner label to a specific Virtual Routing and Forwarding (VRF). This mapping allows examination of the encapsulated IP header at an egress VPN router. For SRX Series Firewalls, the `vrf-table-label` statement is currently supported only on physical interfaces. As a workaround, deactivate `vrf-table-label` or use physical interfaces.

Figure 104: Extranet Configuration of MBGP MVPN with P2MP LDP LSPs as Data Plane



See [Configuring Virtual Loopback Tunnels for VRF Table Lookup](#) for more information.

Overview

IN THIS SECTION

- [Topology | 754](#)

This topic describes how P2MP LDP LSPs can be configured as the data plane for intra-AS selective provider tunnels. Selective P2MP LSPs are triggered only based on the bandwidth threshold of a particular customer's multicast stream. A separate P2MP LDP LSP is set up for a given customer source and customer group pair (C-S, C-G) by a PE router. The C-S is behind the PE router that belongs in the sender site set. Aggregation of intra-AS selective provider tunnels across MVPNs is not supported.

When you configure selective provider tunnels, leaves discover the P2MP LSP root as follows. A PE router with a receiver for a customer multicast stream behind it needs to discover the identity of the PE router (and the provider tunnel information) with the source of the customer multicast stream behind it. This information is auto-discovered dynamically using the S-PMSI AD routes originated by the PE router with the C-S behind it.

The Junos OS also supports P2MP LDP LSPs as the data plane for intra-AS inclusive provider tunnels. These tunnels are triggered based on the MVPN configuration. A separate P2MP LSP LSP is set up for a


given MVPN by a PE router that belongs in the sender site set. This PE router is the root of the P2MP LSP. Aggregation of intra-AS inclusive provider tunnels across MVPNs is not supported.

When you configure inclusive provider tunnels, leaves discover the P2MP LSP root as follows. A PE router with a receiver site for a given MVPN needs to discover the identities of PE routers (and the provider tunnel information) with sender sites for that MVPN. This information is auto-discovered dynamically using the intra-AS auto-discovery routes originated by the PE routers with sender sites.

Topology

Figure 105 on page 754 shows the topology used in this example.

Figure 105: P2MP LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs



In Figure 105 on page 754, the routers perform the following functions:

- R1 and R2 are provider (P) routers.

- R0, R3, R4, and R5 are provider edge (PE) routers.
- MBGP MVPN is configured on all PE routers.
- Two VPNs are defined: green and red.
- Router R0 serves both green and red CE routers in separate routing instances.
- Router R3 is connected to a green CE router.
- Router R5 is connected to overlapping green and red CE routers in a single routing instance.
- Router R4 is connected to overlapping green and red CE routers in a single routing instance.
- OSPF and multipoint LDP (mLDP) are running in the core.
- Router R1 is a route reflector (RR), and router R2 is a redundant RR.
- Routers R0, R3, R4, and R5 are client internal BGP (IBGP) peers.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 755](#)
- [Procedure | 756](#)
- [Results | 757](#)

CLI Quick Configuration

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

```
set protocols ldp interface fe-0/2/1.0
set protocols ldp interface fe-0/2/3.0
set protocols ldp p2mp
set routing-instance red instance-type vrf
set routing-instance red interface vt-0/1/0.1
set routing-instance red interface lo0.1
set routing-instance red route-distinguisher 10.254.1.1:1
```

```
set routing-instance red provider-tunnel ldp-p2mp
set routing-instance red provider-tunnel selective group 224.1.1.1/32 source 192.168.1.1/32 ldp-
p2mp
```

Procedure

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 [Junos OS CLI User Guide](#).

To configure P2MP LDP LSPs as the data plane for intra-AS MBGP MVPNs:

1. Configure LDP on all routers.

```
[edit protocols ldp]
user@host# set interface fe-0/2/1.0
user@host# set interface fe-0/2/3.0
user@host# set p2mp
```

2. Configure the provider tunnel.

```
[edit routing-instance red ]
user@host# set instance-type vrf
user@host# set interface vt-0/1/0.1
user@host# set interface lo0.1
user@host# set route-distinguisher 10.254.1.1:1
user@host# set provider-tunnel ldp-p2mp
```

3. Configure the selective provider tunnel.

```
user@host# set provider-tunnel selective group 224.1.1.1/32 source 192.168.1.1/32 ldp-p2mp
```

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

```
user@host# commit
```

Results

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

```
user@host# show protocols
ldp {
  interface fe-0/2/1.0;
  interface fe-0/2/3.0;
  p2mp;
}
```

```
user@host# show routing-instances
red {
  instance-type vrf;
  interface vt-0/1/0.1;
  interface lo0.1;
  route-distinguisher 10.254.1.1:1;
  provider-tunnel {
    ldp-p2mp;
  }
  selective {
    group 224.1.1.1/32 {
      source 192.168.1.1/32 {
        ldp-p2mp;
      }
    }
  }
}
```

Verification

To verify the configuration, run the following commands:

- **ping mpls ldp p2mp** to ping the end points of a P2MP LSP.
- **show ldp database** to display LDP P2MP label bindings and to ensure that the LDP P2MP LSP is signaled.

- **show ldp session detail** to display the LDP capabilities exchanged with the peer. The **Capabilities advertised** and **Capabilities received** fields should include **p2mp**.
- **show ldp traffic-statistics p2mp** to display the data traffic statistics for the P2MP LSP.
- **show mvpn instance**, **show mvpn neighbor**, and **show mvpn c-multicast** to display multicast VPN routing instance information and to ensure that the LDP P2MP LSP is associated with the MVPN as the S-PMSI.
- **show multicast route instance detail** on PE routers to ensure that traffic is received by all the hosts and to display statistics on the receivers.
- **show route label *label* detail** to display the P2MP forwarding equivalence class (FEC) if the label is an input label for an LDP P2MP LSP.

SEE ALSO

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

[Point-to-Multipoint LSPs Overview](#)

Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs

IN THIS SECTION

- [Requirements | 758](#)
- [Overview | 759](#)
- [Configuration | 761](#)
- [Verification | 766](#)

Requirements

The routers used in this example are Juniper Networks M Series Multiservice Edge Routers, T Series Core Routers, or MX Series 5G Universal Routing Platforms. When using ingress replication for IP multicast, each participating router must be configured with BGP 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, depending on the configuration of the provider tunnel in the routing instance.

Overview

IN THIS SECTION

- [Topology](#) | 759

The ingress-replication provider tunnel type uses unicast tunnels between routers to create a multicast distribution tree.

The [mpls-internet-multicast](#) routing instance type uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP (or Next Gen) MVPN. Ingress replication can also be configured when using MVPN to carry multicast data between PE routers.

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 IP multicast are associated only with the default routing instance (`inet.0`), not with a configured routing instance. The [mpls-internet-multicast](#) routing instance type is configured for the default master instance on each router, and is also included at the `[edit protocols pim]` hierarchy level in the default instance.

For each [mpls-internet-multicast](#) routing instance, the `ingress-replication` statement is required under the `provider-tunnel` statement and also under the `[edit routing-instances routing-instance-name provider-tunnel selective group source]` hierarchy level.

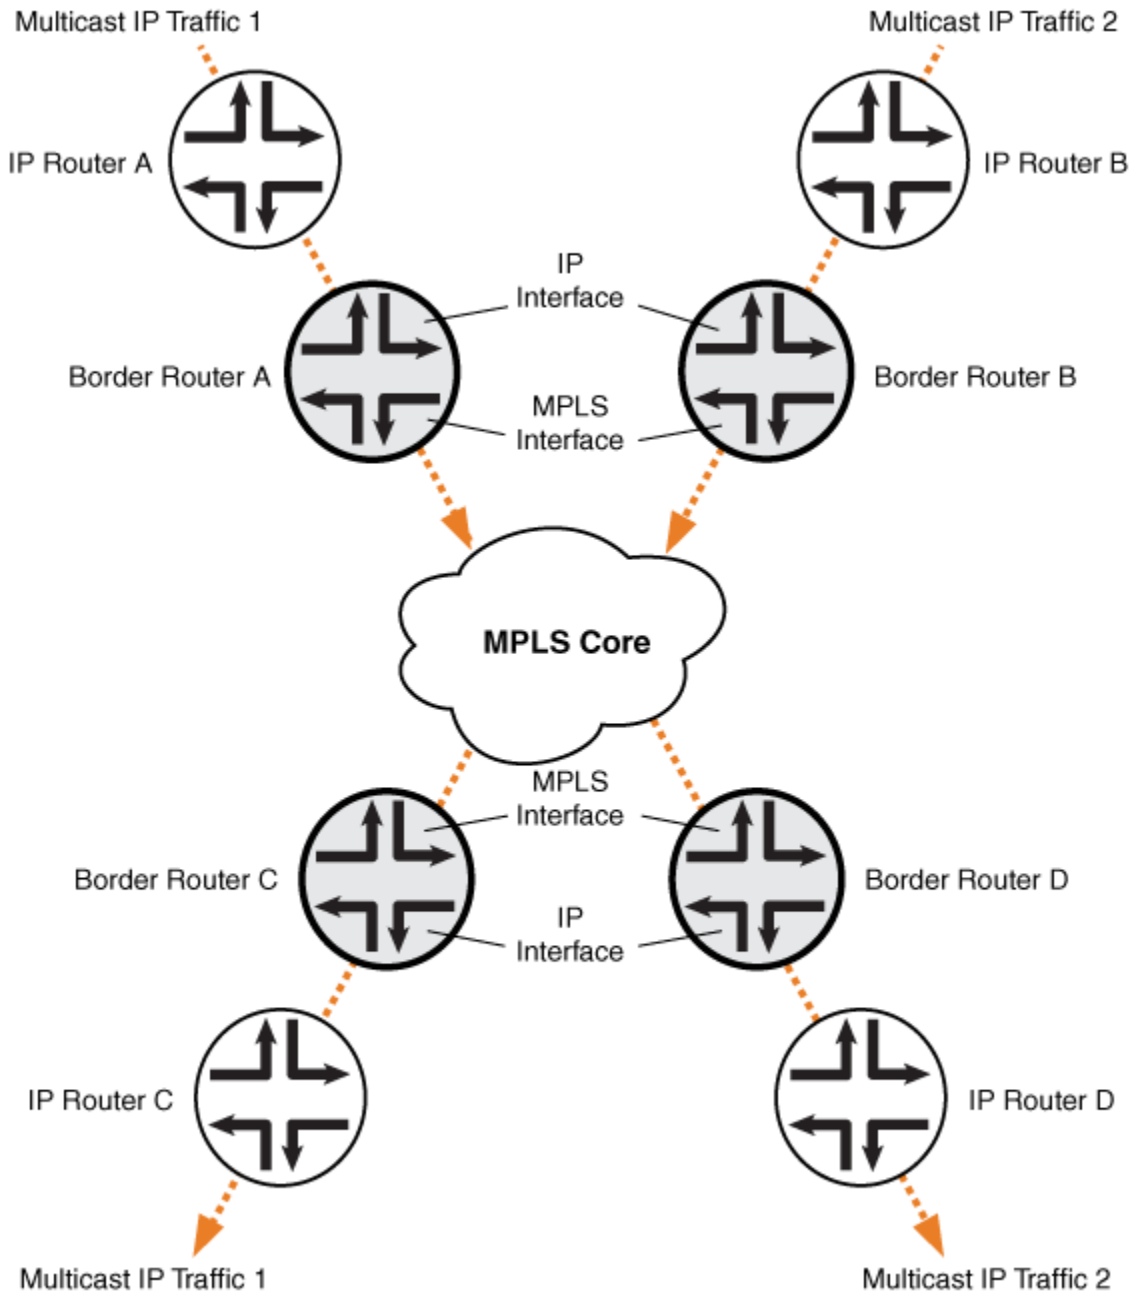
When a new destination needs to be added to the ingress replication provider tunnel, the resulting behavior differs depending on how the ingress replication provider tunnel is configured:

- [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.

Topology

The IP topology consists of routers on the edge of the IP multicast domain. Each router has a set of IP interfaces configured toward the MPLS cloud and a set of interfaces configured toward the IP routers. See [Figure 106 on page 760](#). 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 IBGP session for the control plane.

Figure 106: Internet Multicast Topology



g040632

Configuration

IN THIS SECTION

- [Procedure | 761](#)

Procedure

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.

Border Router C

```
set protocols mpls ipv6-tunneling
set protocols mpls interface all
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.10.61
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet6 unicast
set protocols bgp group ibgp family inet6-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp family inet6-mvpn signaling
set protocols bgp group ibgp export to-bgp
set protocols bgp group ibgp neighbor 10.255.10.97
set protocols bgp group ibgp neighbor 10.255.10.55
set protocols bgp group ibgp neighbor 10.255.10.57
set protocols bgp group ibgp neighbor 10.255.10.59
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 lo0.0
set protocols ospf area 0.0.0.0 interface so-1/3/1.0
set protocols ospf area 0.0.0.0 interface so-0/3/0.0
set protocols ospf3 area 0.0.0.0 interface lo0.0
set protocols ospf3 area 0.0.0.0 interface so-1/3/1.0
set protocols ospf3 area 0.0.0.0 interface so-0/3/0.0
```

```

set protocols ldp interface all
set protocols pim rp static address 192.0.2.2
set protocols pim rp static address 2::192.0.2.2
set protocols pim interface fe-0/1/0.0
set protocols pim mpls-internet-multicast
set routing-instances test instance-type mpls-internet-multicast
set routing-instances test provider-tunnel ingress-replication label-switched-path
set routing-instances test protocols mvpn

```

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](#).

The following example shows how to configure ingress replication on an IP multicast instance with the routing instance type `mpls-internet-multicast`. Additionally, this example shows how to configure a selective provider tunnel that selects a new unicast tunnel each time a new destination needs to be added to the multicast distribution tree.

This example shows the configuration of the link between Border Router C and edge IP Router C, from which Border Router C receives PIM join messages.

1. Enable MPLS.

```

[edit protocols mpls]
user@Border_Router_C# set ipv6-tunneling
user@Border_Router_C# set interface all

```

2. Configure a signaling protocol, such as RSVP or LDP.

```

[edit protocols ldp]
user@Border_Router_C# set interface all

```

3. Configure a full-mesh of IBGP peering sessions.

```

[edit protocols bgp group ibgp]
user@Border_Router_C# set type internal
user@Border_Router_C# set local-address 10.255.10.61
user@Border_Router_C# set neighbor 10.255.10.97

```

```

user@Border_Router_C# set neighbor 10.255.10.55
user@Border_Router_C# set neighbor 10.255.10.57
user@Border_Router_C# set neighbor 10.255.10.59
user@Border_Router_C# set export to-bgp

```

4. Configure the multiprotocol BGP-related settings so that the BGP sessions carry the necessary NLRI.

```

[edit protocols bgp group ibgp]
user@Border_Router_C# set family inet unicast
user@Border_Router_C# set family inet-vpn any
user@Border_Router_C# set family inet6 unicast
user@Border_Router_C# set family inet6-vpn any
user@Border_Router_C# set family inet-mvpn signaling
user@Border_Router_C# set family inet6-mvpn signaling

```

5. Configure an interior gateway protocol (IGP).

This example shows a dual stacking configuration with OSPF and OSPF version 3 configured on the interfaces.

```

[edit protocols ospf3]
user@Border_Router_C# set area 0.0.0.0 interface lo0.0
user@Border_Router_C# set area 0.0.0.0 interface so-1/3/1.0
user@Border_Router_C# set area 0.0.0.0 interface so-0/3/0.0
[edit protocols ospf]
user@Border_Router_C# set traffic-engineering
user@Border_Router_C# set area 0.0.0.0 interface fxp0.0 disable
user@Border_Router_C# set area 0.0.0.0 interface lo0.0
user@Border_Router_C# set area 0.0.0.0 interface so-1/3/1.0
user@Border_Router_C# set area 0.0.0.0 interface so-0/3/0.0

```

6. Configure a global PIM instance on the interface facing the edge device.

PIM is not configured in the core.

```

[edit protocols pim]
user@Border_Router_C# set rp static address 192.0.2.2
user@Border_Router_C# set rp static address 2::192.0.2.2
user@Border_Router_C# set interface fe-0/1/0.0
user@Border_Router_C# set mpls-internet-multicast

```

7. Configure the ingress replication provider tunnel to create a new unicast tunnel each time a destination needs to be added to the multicast distribution tree.

```
[edit routing-instances test]
user@Border_Router_C# set instance-type mpls-internet-multicast
user@Border_Router_C# set provider-tunnel ingress-replication label-switched-path
user@Border_Router_C# set protocols mvpn
```



NOTE: Alternatively, use the **label-switched-path-template** statement to configure a point-to-point LSP for the ingress tunnel.

Configure the point-to-point LSP to use the default template settings (this is needed only when using RSVP tunnels). For example:

```
[edit routing-instances test provider-tunnel]
user@Border_Router_C# set ingress-replication label-switched-path label-switched-
path-template default-template
user@Border_Router_C# set selective group 203.0.113.0/24 source 192.168.195.145/32
ingress-replication label-switched-path
```

8. Commit the configuration.

```
user@Border_Router_C# commit
```

Results

From configuration mode, confirm your configuration by issuing the `show protocols` and `show routing-instances` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@Border_Router_C# show protocols
mpls {
  ipv6-tunneling;
  interface all;
}
bgp {
  group ibgp {
    type internal;
```

```
    local-address 10.255.10.61;
    family inet {
        unicast;
    }
    family inet-vpn {
        any;
    }
    family inet6 {
        unicast;
    }
    family inet6-vpn {
        any;
    }
    family inet-mvpn {
        signaling;
    }
    family inet6-mvpn {
        signaling;
    }
    export to-bgp; ## 'to-bgp' is not defined
    neighbor 10.255.10.97;
    neighbor 10.255.10.55;
    neighbor 10.255.10.57;
    neighbor 10.255.10.59;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
        interface so-1/3/1.0;
        interface so-0/3/0.0;
    }
}
ospf3 {
    area 0.0.0.0 {
        interface lo0.0;
        interface so-1/3/1.0;
        interface so-0/3/0.0;
    }
}
```

```

}
ldp {
  interface all;
}
pim {
  rp {
    static {
      address 192.0.2.2;
      address 2::192.0.2.2;
    }
  }
  interface fe-0/1/0.0;
  mpls-internet-multicast;
}

```

```

user@Border_Router_C# show routing-instances
test {
  instance-type mpls-internet-multicast;
  provider-tunnel {
    ingress-replication {
      label-switched-path;
    }
  }
  protocols {
    mvpn;
  }
}

```

Verification

IN THIS SECTION

- [Checking the Ingress Replication Status on Border Router C | 767](#)
- [Checking the Routing Table for the MVPN Routing Instance on Border Router C | 767](#)
- [Checking the MVPN Neighbors on Border Router C | 769](#)
- [Checking the PIM Join Status on Border Router C | 770](#)
- [Checking the Multicast Route Status on Border Router C | 771](#)

- [Checking the Ingress Replication Status on Border Router B | 771](#)
- [Checking the Routing Table for the MVPN Routing Instance on Border Router B | 772](#)
- [Checking the MVPN Neighbors on Border Router B | 773](#)
- [Checking the PIM Join Status on Border Router B | 774](#)
- [Checking the Multicast Route Status on Border Router B | 775](#)

Confirm that the configuration is working properly. The following operational output is for LDP ingress replication SPT-only mode. The multicast source behind IP Router B. The multicast receiver is behind IP Router C.

Checking the Ingress Replication Status on Border Router C

Purpose

Use the `show ingress-replication mvpn` command to check the ingress replication status.

Action

```
user@Border_Router_C> show ingress-replication mvpn

Ingress Tunnel: mvpn:1
Application: MVPN
Unicast tunnels
  Leaf Address      Tunnel-type    Mode      State
  10.255.10.61     P2P LSP       Existing  Up
```

Meaning

The ingress replication is using a point-to-point LSP, and is in the Up state.

Checking the Routing Table for the MVPN Routing Instance on Border Router C

Purpose

Use the `show route table` command to check the route status.

Action

```

user@Border_Router_C> show route table test.mvpn

test.mvpn.0: 5 destinations, 7 routes (5 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:0:0:10.255.10.61/240
    *[BGP/170] 00:45:55, localpref 100, from 10.255.10.61
    AS path: I, validation-state: unverified
    > via so-2/0/1.0
1:0:0:10.255.10.97/240
    *[MVPN/70] 00:47:19, metric2 1
    Indirect
5:0:0:32:192.168.195.106:32:198.51.100.1/240
    *[PIM/105] 00:06:35
    Multicast (IPv4) Composite
    [BGP/170] 00:06:35, localpref 100, from 10.255.10.61
    AS path: I, validation-state: unverified
    > via so-2/0/1.0
6:0:0:1000:32:192.0.2.2:32:198.51.100.1/240
    *[PIM/105] 00:07:03
    Multicast (IPv4) Composite
7:0:0:1000:32:192.168.195.106:32:198.51.100.1/240
    *[MVPN/70] 00:06:35, metric2 1
    Multicast (IPv4) Composite
    [PIM/105] 00:05:35
    Multicast (IPv4) Composite

test.mvpn-inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:0:0:10.255.10.61/432
    *[BGP/170] 00:45:55, localpref 100, from 10.255.10.61
    AS path: I, validation-state: unverified
    > via so-2/0/1.0
1:0:0:10.255.10.97/432
    *[MVPN/70] 00:47:19, metric2 1
    Indirect

```

Meaning

The expected routes are populating the test.mvpn routing table.

Checking the MVPN Neighbors on Border Router C

Purpose

Use the `show mvpn neighbor` command to check the neighbor status.

Action

```
user@Border_Router_C> show mvpn neighbor
```

```
MVPN instance:
```

```
Legend for provider tunnel
```

```
S- Selective provider tunnel
```

```
Legend for c-multicast routes properties (Pr)
```

```
DS -- derived from (*, c-g)      RM -- remote VPN route
```

```
Family : INET
```

```
Instance : test
```

```
  MVPN Mode : SPT-ONLY
```

```
  Neighbor
```

```
  10.255.10.61
```

```
  Inclusive Provider Tunnel
```

```
  INGRESS-REPLICATION:MPLS Label 16:10.255.10.61
```

```
MVPN instance:
```

```
Legend for provider tunnel
```

```
S- Selective provider tunnel
```

```
Legend for c-multicast routes properties (Pr)
```

```
DS -- derived from (*, c-g)      RM -- remote VPN route
```

```
Family : INET6
```

```
Instance : test
```

```
  MVPN Mode : SPT-ONLY
```

```
  Neighbor
```

```
  10.255.10.61
```

```
  Inclusive Provider Tunnel
```

```
  INGRESS-REPLICATION:MPLS Label 16:10.255.10.61
```

Checking the PIM Join Status on Border Router C

Purpose

Use the `show pim join extensive` command to check the PIM join status.

Action

```
user@Border_Router_C> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 198.51.100.1
  Source: *
  RP: 192.0.2.2
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:07:49
  Downstream neighbors:
    Interface: ge-3/0/6.0
      192.0.2.2 State: Join Flags: SRW Timeout: Infinity
      Uptime: 00:07:49 Time since last Join: 00:07:49
  Number of downstream interfaces: 1

Group: 198.51.100.1
  Source: 192.168.195.106
  Flags: sparse
  Upstream protocol: BGP
  Upstream interface: Through BGP
  Upstream neighbor: Through MVPN
  Upstream state: Local RP, Join to Source, No Prune to RP
  Keepalive timeout: 69
  Uptime: 00:06:21
  Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

Checking the Multicast Route Status on Border Router C

Purpose

Use the `show multicast route extensive` command to check the multicast route status.

Action

```
user@Border_Router_C> show multicast route extensive
Instance: master Family: INET

Group: 198.51.100.1
  Source: 192.168.195.106/32
  Upstream interface: lsi.0
  Downstream interface list:
    ge-3/0/6.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 18 kBps, 200 pps, 88907 packets
  Next-hop ID: 1048577
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:07:25

Instance: master Family: INET6
```

Checking the Ingress Replication Status on Border Router B

Purpose

Use the `show ingress-replication mvpn` command to check the ingress replication status.

Action

```
user@Border_Router_B> show ingress-replication mvpn
```

```
Ingress Tunnel: mvpn:1
Application: MVPN
Unicast tunnels
  Leaf Address      Tunnel-type  Mode      State
  10.255.10.97     P2P LSP    Existing  Up
```

Meaning

The ingress replication is using a point-to-point LSP, and is in the Up state.

Checking the Routing Table for the MVPN Routing Instance on Border Router B

Purpose

Use the `show route table` command to check the route status.

Action

```
user@Border_Router_B> show route table test.mvpn

test.mvpn.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:0:0:10.255.10.61/240
    *[MVPN/70] 00:49:26, metric2 1
    Indirect
1:0:0:10.255.10.97/240
    *[BGP/170] 00:48:22, localpref 100, from 10.255.10.97
    AS path: I, validation-state: unverified
    > via so-1/3/1.0
5:0:0:32:192.168.195.106:32:198.51.100.1/240
    *[PIM/105] 00:09:02
    Multicast (IPv4) Composite
    [BGP/170] 00:09:02, localpref 100, from 10.255.10.97
    AS path: I, validation-state: unverified
    > via so-1/3/1.0
7:0:0:1000:32:192.168.195.106:32:198.51.100.1/240
    *[PIM/105] 00:09:02
    Multicast (IPv4) Composite
    [BGP/170] 00:09:02, localpref 100, from 10.255.10.97
    AS path: I, validation-state: unverified
```

```

> via so-1/3/1.0

test.mvpn-inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:0:0:10.255.10.61/432
    *[MVPN/70] 00:49:26, metric2 1
    Indirect
1:0:0:10.255.10.97/432
    *[BGP/170] 00:48:22, localpref 100, from 10.255.10.97
    AS path: I, validation-state: unverified
    > via so-1/3/1.0

```

Meaning

The expected routes are populating the test.mvpn routing table.

Checking the MVPN Neighbors on Border Router B

Purpose

Use the `show mvpn neighbor` command to check the neighbor status.

Action

```

user@Border_Router_B> show mvpn neighbor

MVPN instance:
Legend for provider tunnel
S-   Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Family : INET

Instance : test
MVPN Mode : SPT-ONLY
Neighbor           Inclusive Provider Tunnel
10.255.10.97       INGRESS-REPLICATION:MPLS Label 16:10.255.10.97

MVPN instance:

```

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET6

Instance : test

MVPN Mode : SPT-ONLY

Neighbor

10.255.10.97

Inclusive Provider Tunnel

INGRESS-REPLICATION:MPLS Label 16:10.255.10.97

Checking the PIM Join Status on Border Router B

Purpose

Use the `show pim join extensive` command to check the PIM join status.

Action

```

user@Border_Router_B> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 198.51.100.1
  Source: 192.168.195.106
  Flags: sparse,spt
  Upstream interface: fe-0/1/0.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout: 0
  Uptime: 00:09:39
  Downstream neighbors:
    Interface: Pseudo-MVPN
      Uptime: 00:09:39 Time since last Join: 00:09:39
  Number of downstream interfaces: 1

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```


Checking the Multicast Route Status on Border Router B

Purpose

Use the `show multicast route extensive` command to check the multicast route status.

Action

```
user@Border_Router_B> show multicast route extensive
Instance: master Family: INET

Group: 198.51.100.1
  Source: 192.168.195.106/32
  Upstream interface: fe-0/1/0.0
  Downstream interface list:
    so-1/3/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 18 kBps, 200 pps, 116531 packets
  Next-hop ID: 1048580
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:09:43
```

SEE ALSO

Configuring Routing Instances for an MBGP MVPN

mpls-internet-multicast

ingress-replication

create-new-ucast-tunnel

label-switched-path-template (Multicast)

show ingress-replication mvpn

Example: Configuring MBGP Multicast VPNs

IN THIS SECTION

- [Requirements | 776](#)
- [Overview and Topology | 776](#)
- [Configuration | 777](#)

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

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

IN THIS SECTION

- [Topology | 777](#)

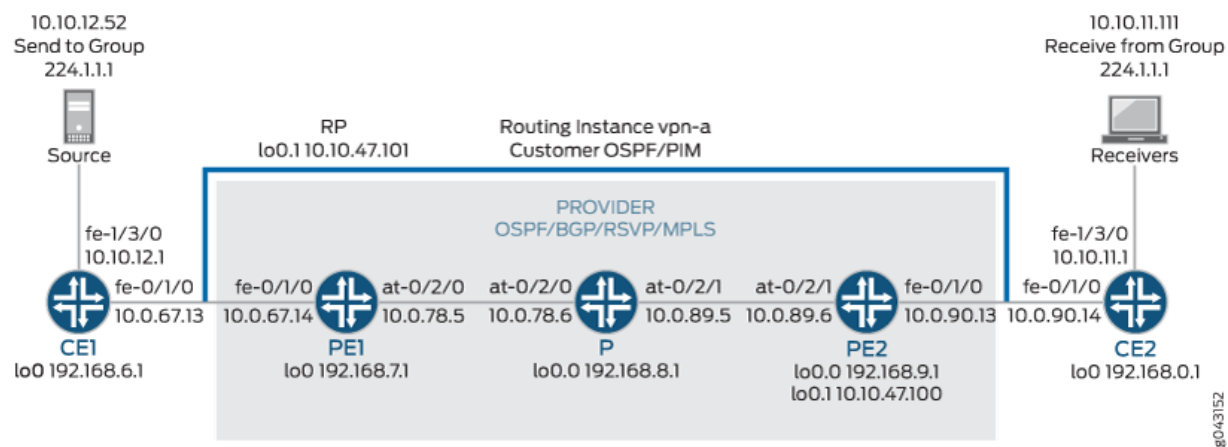
This example shows how to configure the following technologies:

- IPv4
- BGP
- OSPF
- RSVP
- MPLS
- PIM sparse mode
- Static RP

Topology

The topology of the network is shown in [Figure 107 on page 777](#).

Figure 107: Multicast Over Layer 3 VPN Example Topology



Configuration

IN THIS SECTION

- [Configuring Interfaces | 778](#)
- [Configuring OSPF | 780](#)

- [Configuring BGP | 782](#)
- [Configuring RSVP | 783](#)
- [Configuring MPLS | 784](#)
- [Configuring the VRF Routing Instance | 785](#)
- [Configuring PIM | 787](#)
- [Configuring the Provider Tunnel | 787](#)
- [Configuring the Rendezvous Point | 788](#)
- [Results | 789](#)



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 107 on page 777](#), perform the following steps:

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 `100.0` 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 100.0 address. The neighbor addresses are the PE routers' 100.0 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 100.0 address. The neighbor addresses are the 100.0 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 `10.10.47.101` 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 `10.10.47.101` 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 `10.10.47.101` 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:cdb/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;
    }
  }
}
```

```
}
}
```

Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN

IN THIS SECTION

- [Requirements | 800](#)
- [Overview | 800](#)
- [Configuration | 802](#)
- [Verification | 811](#)

This example shows how to configure a PIM-SSM provider tunnel for an MBGP MVPN. The configuration enables service providers to carry customer data in the core. This example shows how to configure PIM-SSM tunnels as inclusive PMSI and uses the unicast routing preference as the metric for determining the single forwarder (instead of the default metric, which is the IP address from the global administrator field in the route-import community).

Requirements

Before you begin:

- Configure the router interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure the BGP-to-OSPF routing policy. See the [Routing Policies, Firewall Filters, and Traffic Policers User Guide](#).

Overview

IN THIS SECTION

- [Topology | 801](#)

When a PE receives a customer join or prune message from a CE, the message identifies a particular multicast flow as belonging either to a source-specific tree (S,G) or to a shared tree (*,G). If the route to the multicast source or RP is across the VPN backbone, then the PE needs to identify the upstream

multicast hop (UMH) for the (S,G) or (*,G) flow. Normally the UMH is determined by the unicast route to the multicast source or RP.

However, in some cases, the CEs might be distributing to the PEs a special set of routes that are to be used exclusively for the purpose of upstream multicast hop selection using the route-import community. More than one route might be eligible, and the PE needs to elect a single forwarder from the eligible UMHs.

The default metric for the single forwarder election is the IP address from the global administrator field in the route-import community. You can configure a router to use the unicast route preference to determine the single forwarder election.

This example includes the following settings.

- **provider-tunnel family inet pim-ssm group-address**—Specifies a valid SSM VPN group address. The SSM VPN group address and the source address are advertised by the type-1 autodiscovery route. On receiving an autodiscovery route with the SSM VPN group address and the source address, a PE router sends an (S,G) join in the provider space to the PE advertising the autodiscovery route. All PE routers exchange their PIM-SSM VPN group address to complete the inclusive provider multicast service interface (I-PMSI). Unlike a PIM-ASM provider tunnel, the PE routers can choose a different VPN group address because the (S,G) joins are sent directly toward the source PE.




NOTE: Similar to a PIM-ASM provider tunnel, PIM must be configured in the default master instance.

- **unicast-umh-election**—Specifies that the PE router uses the unicast route preference to determine the single-forwarder election.

Topology

Figure 108 on page 801 shows the topology used in this example.

Figure 108: PIM-SSM Provider Tunnel for an MBGP MVPN Topology



Configuration

IN THIS SECTION

- [Procedure | 802](#)
- [Results | 806](#)

Procedure

CLI Quick Configuration

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

```
set interfaces fe-0/2/0 unit 0 family inet address 192.168.195.109/30
set interfaces fe-0/2/1 unit 0 family inet address 192.168.195.5/27
set interfaces fe-0/2/2 unit 0 family inet address 10.20.1.1/30
set interfaces fe-0/2/2 unit 0 family iso
set interfaces fe-0/2/2 unit 0 family mpls
set interfaces lo0 unit 1 family inet address 10.10.47.100/32
set interfaces lo0 unit 1 family inet address 192.168.195.1/32 primary
set interfaces lo0 unit 2 family inet address 10.10.48.100/32
set protocols mpls interface all set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-preference 120
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.255.112.155
set protocols isis level 1 disable set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols ospf traffic-engineering
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 pim rp static address 10.255.112.155
set protocols pim interface all mode sparse-dense
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
```

```

set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface fe-0/2/1.0
set routing-instances VPN-A interface lo0.1
set routing-instances VPN-A route-distinguisher 10.255.112.199:100
set routing-instances VPN-A provider-tunnel family inet pim-ssm group-address 233.252.0.1
set routing-instances VPN-A vrf-target target:10:100
set routing-instances VPN-A vrf-table-label
set routing-instances VPN-A routing-options auto-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 lo0.1
set routing-instances VPN-A protocols ospf area 0.0.0.0 interface fe-0/2/1.0
set routing-instances VPN-A protocols pim rp static address 10.10.47.101
set routing-instances VPN-A protocols pim interface lo0.1 mode sparse-dense
set routing-instances VPN-A protocols pim interface lo0.1 version 2
set routing-instances VPN-A protocols pim interface fe-0/2/1.0 mode sparse-dense
set routing-instances VPN-A protocols pim interface fe-0/2/1.0 version 2
set routing-instances VPN-A protocols mvpn unicast-umh-election
set routing-instances VPN-B instance-type vrf
set routing-instances VPN-B interface fe-0/2/0.0
set routing-instances VPN-B interface lo0.2
set routing-instances VPN-B route-distinguisher 10.255.112.199:200
set routing-instances VPN-B provider-tunnel family inet pim-ssm group-address 233.252.0.2
set routing-instances VPN-B vrf-target target:10:200
set routing-instances VPN-B vrf-table-label
set routing-instances VPN-B routing-options auto-export
set routing-instances VPN-B protocols ospf export bgp-to-ospf
set routing-instances VPN-B protocols ospf area 0.0.0.0 interface lo0.2
set routing-instances VPN-B protocols ospf area 0.0.0.0 interface fe-0/2/0.0
set routing-instances VPN-B protocols pim rp static address 10.10.48.101
set routing-instances VPN-B protocols pim interface lo0.2 mode sparse-dense
set routing-instances VPN-B protocols pim interface lo0.2 version 2
set routing-instances VPN-B protocols pim interface fe-0/2/0.0 mode sparse-dense
set routing-instances VPN-B protocols pim interface fe-0/2/0.0 version 2
set routing-instances VPN-B protocols mvpn unicast-umh-election
set routing-options autonomous-system 65100

```

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 [Junos OS CLI User Guide](#).

To configure a PIM-SSM provider tunnel for an MBGP MVPN:

1. Configure the interfaces in the master routing instance on the PE routers. This example shows the interfaces for one PE router.

```
[edit interfaces]
user@host# set fe-0/2/0 unit 0 family inet address 192.168.195.109/30
user@host# set fe-0/2/1 unit 0 family inet address 192.168.195.5/27
user@host# set fe-0/2/2 unit 0 family inet address 10.20.1.1/30
user@host# set fe-0/2/2 unit 0 family iso
user@host# set fe-0/2/2 unit 0 family mpls
user@host# set lo0 unit 1 family inet address 10.10.47.100/32
user@host# set lo0 unit 2 family inet address 10.10.48.100/32
```

2. Configure the autonomous system number in the global routing options. This is required in MBGP MVPNs.

```
[edit routing-options]
user@host# set autonomous-system 65100
```

3. Configure the routing protocols in the master routing instance on the PE routers.

```
[edit protocols bgp group ibgp]
user@host# set type internal
user@host# set family inet-vpn any
user@host# set family inet-mvpn signaling
user@host# set neighbor 10.255.112.155
[edit protocols isis]
user@host# set level 1 disable
user@host# set interface all
user@host# set interface fxp0.0 disable
[edit protocols ospf]
user@host# set traffic-engineering
user@host# set area 0.0.0.0 interface all
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set protocols ldp interface all
[edit protocols pim]
user@host# set rp static address 10.255.112.155
user@host# set interface all mode sparse-dense
```

```

user@host# set interface all version 2
user@host# set interface fxp0.0 disable

```

4. Configure routing instance VPN-A.

```

[edit routing-instances VPN-A]
user@host# set instance-type vrf
user@host# set interface fe-0/2/1.0
user@host# set interface lo0.1
user@host# set route-distinguisher 10.255.112.199:100
user@host# set provider-tunnel family inet pim-ssm group-address 232.252.0.1
user@host# set vrf-target target:10:100
user@host# set vrf-table-label
user@host# set routing-options auto-export
user@host# set protocols ospf export bgp-to-ospf
user@host# set protocols ospf area 0.0.0.0 interface lo0.1
user@host# set protocols ospf area 0.0.0.0 interface fe-0/2/1.0
user@host# set protocols pim rp static address 10.10.47.101
user@host# set protocols pim interface lo0.1 mode sparse-dense
user@host# set protocols pim interface lo0.1 version 2
user@host# set protocols pim interface fe-0/2/1.0 mode sparse-dense
user@host# set protocols pim interface fe-0/2/1.0 version 2
user@host# set protocols mvpn unicast-umh-election

```

5. Configure routing instance VPN-B.

```

[edit routing-instances VPN-B]
user@host# set instance-type vrf
user@host# set interface fe-0/2/0.0
user@host# set interface lo0.2
user@host# set route-distinguisher 10.255.112.199:200
user@host# set provider-tunnel family inet pim-ssm group-address 232.252.0.2
user@host# set vrf-target target:10:200
user@host# set vrf-table-label
user@host# set routing-options auto-export
user@host# set protocols ospf export bgp-to-ospf
user@host# set protocols ospf area 0.0.0.0 interface lo0.2
user@host# set protocols ospf area 0.0.0.0 interface fe-0/2/0.0
user@host# set protocols pim rp static address 10.10.48.101
user@host# set protocols pim interface lo0.2 mode sparse-dense
user@host# set protocols pim interface lo0.2 version 2

```

```

user@host# set protocols pim interface fe-0/2/0.0 mode sparse-dense
user@host# set protocols pim interface fe-0/2/0.0 version 2
user@host# set protocols mvpn unicast-umh-election

```

6. Configure the topology such that the BGP route to the source advertised by PE1 has a higher preference than the BGP route to the source advertised by PE2.

```

[edit protocols bgp]
user@host# set group ibgp local-preference 120

```

7. Configure a higher primary loopback address on PE2 than on PE1. This ensures that PE2 is the MBGP MVPN single-forwarder election winner.

```

[edit]
user@host# set interface lo0 unit 1 family inet address 192.168.195.1/32 primary

```

8. Configure the unicast-umh-election statement on PE3.

```

[edit]
user@host# set routing-instances VPN-A protocols mvpn unicast-umh-election
user@host# set routing-instances VPN-B protocols mvpn unicast-umh-election

```

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

```

user@host# commit

```

Results

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

```

user@host# show interfaces
fe-0/2/0 {
  unit 0 {
    family inet {
      address 192.168.195.109/30;

```



```
    }
  }
}
fe-0/2/1 {
  unit 0 {
    family inet {
      address 192.168.195.5/27;
    }
  }
}
fe-0/2/2 {
  unit 0 {
    family inet {
      address 10.20.1.1/30;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.10.47.100/32;
      address 192.168.195.1/32 {
        primary;
      }
    }
  }
  unit 2 {
    family inet {
      address 10.10.48.100/32;
    }
  }
}
}
```

```
user@host# show protocols
mpls {
  interface all;
}
bgp {
  group ibgp {
```

```
    type internal;
    local-preference 120;
    family inet-vpn {
        any;
    }
    family inet-mvpn {
        signaling;
    }
    neighbor 10.255.112.155;
}
}
isis {
    level 1 disable;
    interface all;
    interface fxp0.0 {
        disable;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}
ldp {
    interface all;
}
}
pim {
    rp {
        static {
            address 10.255.112.155;
        }
    }
}
}
interface all {
    mode sparse-dense;
    version 2;
}
}
interface fxp0.0 {
    disable;
}
```

```
}  
}
```

```
user@host# show routing-instances  
VPN-A {  
  instance-type vrf;  
  interface fe-0/2/1.0;  
  interface lo0.1;  
  route-distinguisher 10.255.112.199:100;  
  provider-tunnel {  
    family inet  
      pim-ssm {  
        group-address 233.252.0.1;  
      }  
  }  
  vrf-target target:10:100;  
  vrf-table-label;  
  routing-options {  
    auto-export;  
  }  
  protocols {  
    ospf {  
      export bgp-to-ospf;  
      area 0.0.0.0 {  
        interface lo0.1;  
        interface fe-0/2/1.0;  
      }  
    }  
    pim {  
      rp {  
        static {  
          address 10.10.47.101;  
        }  
      }  
      interface lo0.1 {  
        mode sparse-dense;  
        version 2;  
      }  
      interface fe-0/2/1.0 {  
        mode sparse-dense;  
        version 2;  
      }  
    }  
  }  
}
```

```
    }
  }
  mvpn {
    unicast-umh-election;
  }
}
VPN-B {
  instance-type vrf;
  interface fe-0/2/0.0;
  interface lo0.2;
  route-distinguisher 10.255.112.199:200;
  provider-tunnel {
    family inet {
      pim-ssm {
        group-address 233.252.0.2;
      }
    }
  }
  vrf-target target:10:200;
  vrf-table-label;
  routing-options {
    auto-export;
  }
  protocols {
    ospf {
      export bgp-to-ospf;
      area 0.0.0.0 {
        interface lo0.2;
        interface fe-0/2/0.0;
      }
    }
  }
  pim {
    rp {
      static {
        address 10.10.48.101;
      }
    }
  }
  interface lo0.2 {
    mode sparse-dense;
    version 2;
  }
  interface fe-0/2/0.0 {
    mode sparse-dense;
  }
}
```

```

        version 2;
    }
}
mvpn {
    unicast-umh-election;
}
}
}

```

```

fe-0/2/0 {
    unit 0 {
        family inet {
            address 192.168.195.109/30;
        }
    }
}
fe-0/2/1 {
    unit 0 {
        family inet {
            address 192.168.195.5/27;
        }
    }
}
}

```

```

user@host# show routing-options
autonomous-system 65100;

```

Verification

To verify the configuration, start the receivers and the source. PE3 should create type-7 customer multicast routes from the local joins. Verify the source-tree customer multicast entries on all PE routers. PE3 should choose PE1 as the upstream PE toward the source. PE1 receives the customer multicast route from the egress PEs and forwards data on the PSMI to PE3.

To confirm the configuration, run the following commands:

- **show route table VPN-A.mvpn.0 extensive**
- **show multicast route extensive instance VPN-A**

SEE ALSO

[Configuring a Selective Provider Tunnel Using Wildcards | 1014](#)

Configuring PIM Provider Tunnels for an MBGP MVPN

Example: Allowing MBGP MVPN Remote Sources

IN THIS SECTION

- [Requirements | 812](#)
- [Overview | 812](#)
- [Configuration | 814](#)
- [Verification | 818](#)

This example shows how to configure an MBGP MVPN that allows remote sources, even when there is no PIM neighborship toward the upstream router.

Requirements

Before you begin:

- Configure the router interfaces. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure the point-to-multipoint static LSP. See [Configuring Point-to-Multipoint LSPs for an MBGP MVPN](#).

Overview

IN THIS SECTION

- [Topology | 813](#)

In this example, a remote CE router is the multicast source. In an MBGP MVPN, a PE router has the PIM interface hello interval set to zero, thereby creating no PIM neighborship. The PIM upstream state is

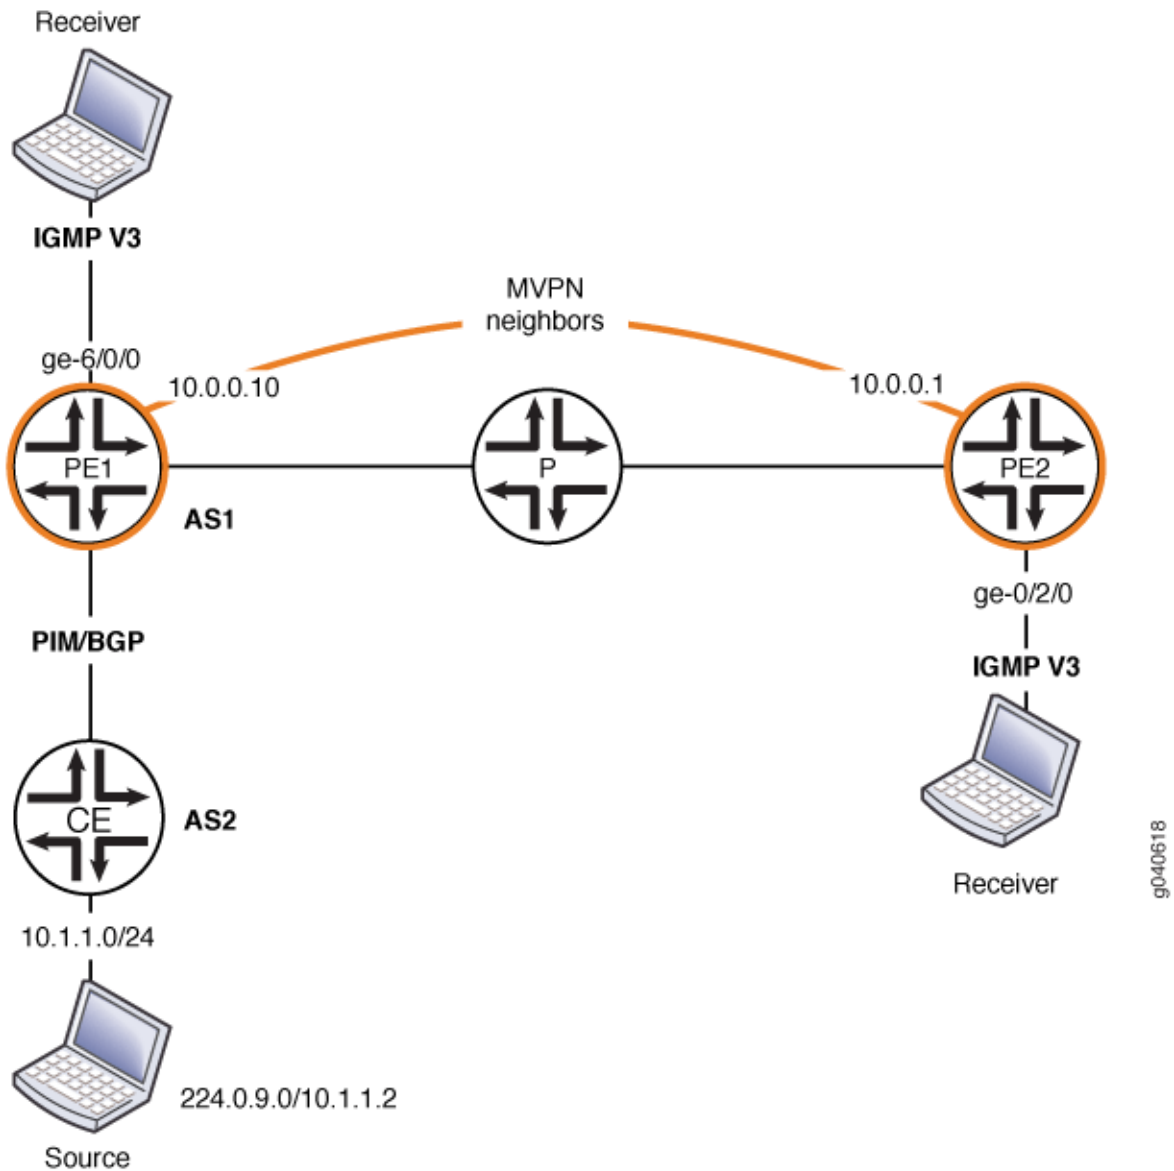
None. In this scenario, directly connected receivers receive traffic in the MBGP MVPN only if you configure the ingress PE's upstream logical interface to accept remote sources. If you do not configure the ingress PE's logical interface to accept remote sources, the multicast route is deleted and the local receivers are no longer attached to the flood next hop.

This example shows the configuration on the ingress PE router. A static LSP is used to receive traffic from the remote source.

Topology

[Figure 109 on page 814](#) shows the topology used in this example.

Figure 109: MBGP MVPN Remote Source



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 815](#)
- [Procedure | 815](#)
- [Results | 817](#)

CLI Quick Configuration

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

```

set routing-instances vpn-A instance-type vrf
set routing-instances vpn-A interface ge-1/0/0.213
set routing-instances vpn-A interface ge-1/0/0.484
set routing-instances vpn-A interface ge-1/0/1.200
set routing-instances vpn-A interface ge-1/0/2.0
set routing-instances vpn-A interface ge-1/0/7.0
set routing-instances vpn-A interface vt-1/1/0.0
set routing-instances vpn-A route-distinguisher 10.0.0.10:04
set routing-instances vpn-A provider-tunnel rsvp-te label-switched-path-template mvpn-dynamic
set routing-instances vpn-A provider-tunnel selective group 224.0.9.0/32 source 10.1.1.2/32 rsvp-
te static-lsp mvpn-static
set routing-instances vpn-A vrf-target target:65000:04
set routing-instances vpn-A protocols bgp group 1a type external
set routing-instances vpn-A protocols bgp group 1a peer-as 65213
set routing-instances vpn-A protocols bgp group 1a neighbor 10.2.213.9
set routing-instances vpn-A protocols pim interface all hello-interval 0
set routing-instances vpn-A protocols pim interface ge-1/0/2.0 accept-remote-source
set routing-instances vpn-A protocols mvpn
set routing-options autonomous-system 100

```

Procedure

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 [Junos OS CLI User Guide](#).

To allow remote sources:

1. On the ingress PE router, configure the interfaces in the routing instance.

```

[edit routing-instances vpn-A]
user@host# set instance-type vrf
user@host# set interface ge-1/0/0.213

```

```

user@host# set interface ge-1/0/0.484
user@host# set interface ge-1/0/1.200
user@host# set interface ge-1/0/2.0
user@host# set interface ge-1/0/7.0
user@host# set interface vt-1/1/0.0

```

2. Configure the autonomous system number in the global routing options. This is required in MBGP MVPNs.

```

user@host# set routing-options autonomous-system 100

```

3. Configure the route distinguisher and the VRF target.

```

[edit routing-instances vpn-A]
user@host# set route-distinguisher 10.0.0.10:04
user@host# set vrf-target target:65000:04

```

4. Configure the provider tunnel.

```

[edit routing-instances vpn-A]
user@host# set provider-tunnel rsvp-te label-switched-path-template mvpn-dynamic
user@host# set provider-tunnel selective group 224.0.9.0/32 source 10.1.1.2/32 rsvp-te static-lsp mvpn-static

```

5. Configure BGP in the routing instance.

```

[edit routing-instances vpn-A]
user@host# set protocols bgp group 1a type external
user@host# set protocols bgp group 1a peer-as 65213
user@host# set protocols bgp group 1a neighbor 10.2.213.9

```

6. Configure PIM in the routing instance, including the accept-remote-source statement on the incoming logical interface.

```

[edit routing-instances vpn-A]
user@host# set protocols pim interface all hello-interval 0
user@host# set protocols pim interface ge-1/0/2.0 accept-remote-source

```

7. Enable the MVPN Protocol in the routing instance.

```
[edit routing-instances vpn-A]
user@host# set protocols mvpn
```

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

```
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the `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@host# show routing-instances
routing-instances {
  vpn-A {
    instance-type vrf;
    interface ge-1/0/0.213;
    interface ge-1/0/0.484;
    interface ge-1/0/1.200;
    interface vt-1/1/0.0;
    interface ge-1/0/2.0;
    interface ge-1/0/7.0;
    route-distinguisher 10.0.0.10:04;
    provider-tunnel {
      rsvp-te {
        label-switched-path-template {
          mvpn-dynamic;
        }
      }
      selective {
        group 224.0.9.0/32 {
          source 10.1.1.2/32 {
            rsvp-te {
              static-lsp mvpn-static;
            }
          }
        }
      }
    }
  }
}
```

```

    }
  }
  vrf-target target:65000:04;
  protocols {
    bgp {
      group 1a {
        type external;
        peer-as 65213;
        neighbor 10.2.213.9;
      }
    }
    pim {
      interface all {
        hello-interval 0;
      }
      interface ge-1/0/2.0 {
        accept-remote-source;
      }
    }
    mvpn;
  }
}

```

```

user@host# show routing-options
autonomous-system 100;

```

Verification

To verify the configuration, run the following commands:

- **show mpls lsp p2mp**
- **show multicast route instance vpn-A extensive**
- **show mvpn c-multicast**
- **show pim join instance vpn-A extensive**
- **show route forwarding-table destination *destination***
- **show route table vpn-A.mvpn.0**

SEE ALSO

[Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN | 800](#)

accept-remote-source

Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

IN THIS SECTION

- [Requirements | 819](#)
- [Overview | 819](#)
- [Configuration | 820](#)
- [Verification | 832](#)

This example shows how to configure an multiprotocol BGP multicast VPN (also called Next-Generation MVPN) with BGP route flap damping.

Requirements

This example uses Junos OS Release 12.2. BGP route flap damping support for MBGP MVPN, specifically, and on an address family basis, in general, is introduced in Junos OS Release 12.2.

Overview

IN THIS SECTION

- [Topology | 820](#)


BGP route flap damping helps to diminish route instability caused by routes being repeatedly withdrawn and readvertised when a link is intermittently failing.

This example uses the default damping parameters and demonstrates an MBGP MVPN scenario with three provider edge (PE) routing devices, three customer edge (CE) routing devices, and one provider (P) routing device.

Topology

Figure 110 on page 820 shows the topology used in this example.

Figure 110: MBGP MVPN with BGP Route Flap Damping



On PE Device R4, BGP route flap damping is configured for address family `inet-mvpn`. A routing policy called `dampPolicy` uses the `nLri-route-type` match condition to damp only MVPN route types 3, 4, and 5. All other MVPN route types are not damped.

This example shows the full configuration on all devices in the "CLI Quick Configuration" on page 821 section. The "Configuring Device R4" on page 825 section shows the step-by-step configuration for PE Device R4.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 821](#)
- [Configuring Device R4 | 825](#)
- [Results | 828](#)

CLI Quick Configuration

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

Device R1

```
set interfaces 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 172.16.1.1/32
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 172.16.100.1
set protocols pim interface all
set routing-options router-id 172.16.1.1
```

Device R2

```
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 172.16.1.2/32
set interfaces lo0 unit 102 family inet address 172.16.100.1/32
set protocols mpls interface ge-1/2/1.5
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.2
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 172.16.1.4
set protocols bgp group ibgp neighbor 172.16.1.5
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 rp static address 172.16.1.2 with 172.16.4.1100.1
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 172.16.1.2
set routing-options autonomous-system 1001

```

Device R3

```

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
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 172.16.1.3/32
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 172.16.1.3

```

Device R4

```

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 172.16.1.4/32
set interfaces lo0 unit 104 family inet address 172.16.100.1/32
set protocols rsvp interface all aggregate
set protocols mpls interface all
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.4
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling damping
set protocols bgp group ibgp neighbor 172.16.1.2 import dampPolicy
set protocols bgp group ibgp neighbor 172.16.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
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 dampPolicy term term1 from family inet-mvpn
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 3
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 4
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 5
set policy-options policy-statement dampPolicy term term1 then accept
set policy-options policy-statement dampPolicy then damping no-damp
set policy-options policy-statement dampPolicy then accept
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set policy-options damping no-damp disable
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 static address 172.16.100.1
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 172.16.1.4
set routing-options autonomous-system 64501
```

Device R5

```
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 172.16.1.5/32
set interfaces lo0 unit 105 family inet address 172.16.100.5/32
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.4
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 172.16.100.2
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 172.16.1.5
set routing-options autonomous-system 1001
```

Device R6

```

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 172.16.1.6/32
set protocols sap listen 233.1.1.1
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 172.16.100.2
set protocols pim interface all
set routing-options router-id 172.16.1.6

```

Device R7

```

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 172.16.1.7/32
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 172.16.100.2
set protocols pim interface all
set routing-options router-id 172.16.1.7

```

Configuring Device R4

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 [Junos OS CLI User Guide](#).

To configure Device R4:

1. Configure the interfaces.

```

[edit interfaces]
user@R4# set ge-1/2/0 unit 10 family inet address 10.1.1.10/30
user@R4# set ge-1/2/0 unit 10 family mpls
user@R4# set ge-1/2/1 unit 17 family inet address 10.1.1.17/30
user@R4# set ge-1/2/1 unit 17 family mpls

```

```

user@R4# set vt-1/2/0 unit 4 family inet
user@R4# set lo0 unit 4 family inet address 172.16.1.4/32
user@R4# set lo0 unit 104 family inet address 172.16.100.4/32

```

2. Configure MPLS and the signaling protocols on the interfaces.

```

[edit protocols]
user@R4# set mpls interface all
user@R4# set mpls interface ge-1/2/0.10
user@R4# set rsvp interface all aggregate
user@R4# set ldp interface ge-1/2/0.10
user@R4# set ldp p2mp

```

3. Configure BGP.

The BGP configuration enables BGP route flap damping for the `inet-mvpn` address family. The BGP configuration also imports into the routing table the routing policy called `dampPolicy`. This policy is applied to neighbor PE Device R2.

```

[edit protocols bgp group ibgp]
user@R4# set type internal
user@R4# set local-address 172.16.1.4
user@R4# set family inet-vpn unicast
user@R4# set family inet-vpn any
user@R4# set family inet-mvpn signaling damping
user@R4# set neighbor 172.16.1.2 import dampPolicy
user@R4# set neighbor 172.16.1.5

```

4. Configure an interior gateway protocol.

```

[edit protocols ospf]
user@R4# set traffic-engineering
[edit protocols ospf area 0.0.0.0]
user@R4# set interface all
user@R4# set interface lo0.4 passive
user@R4# set interface ge-1/2/0.10

```

5. Configure a damping policy that uses the `nlri-route-type` match condition to damp only MVPN route types 3, 4, and 5.

```
[edit policy-options policy-statement dampPolicy term term1]
user@R4# set from family inet-mvpn
user@R4# set from nlri-route-type 3
user@R4# set from nlri-route-type 4
user@R4# set from nlri-route-type 5
user@R4# set then accept
```

6. Configure the damping policy to disable BGP route flap damping.

The `no-damp` policy (`damping no-damp disable`) causes any damping state that is present in the routing table to be deleted. The `then damping no-damp` statement applies the `no-damp` policy as an action and has no `from` match conditions. Therefore, all routes that are not matched by `term1` are matched by this term, with the result that all other MVPN route types are not damped.

```
[edit policy-options policy-statement dampPolicy]
user@R4# set then damping no-damp
user@R4# set then accept
[edit policy-options]
user@R4# set damping no-damp disable
```

7. Configure the `parent_vpn_routes` to accept all other BGP routes that are not from the `inet-mvpn` address family.

This policy is applied as an OSPF export policy in the routing instance.

```
[edit policy-options policy-statement parent_vpn_routes]
user@R4# set from protocol bgp
user@R4# set then accept
```

8. Configure the VPN routing and forwarding (VRF) instance.

```
[edit routing-instances vpn-1]
user@R4# set instance-type vrf
user@R4# set interface vt-1/2/0.4
user@R4# set interface ge-1/2/1.17
user@R4# set interface lo0.104
user@R4# set route-distinguisher 100:100
```

```

user@R4# set vrf-target target:1:1
user@R4# set protocols ospf export parent_vpn_routes
user@R4# set protocols ospf area 0.0.0.0 interface lo0.104 passive
user@R4# set protocols ospf area 0.0.0.0 interface ge-1/2/1.17
user@R4# set protocols pim rp static address 172.16.100.2
user@R4# set protocols pim interface ge-1/2/1.17 mode sparse
user@R4# set protocols mvpn

```

9. Configure the router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@R4# set router-id 172.16.1.4
user@R4# set autonomous-system 1001

```

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

```

user@R4# commit

```

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@R4# show interfaces
ge-1/2/0 {
  unit 10 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 17 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}

```

```
    }  
  }  
  vt-1/2/0 {  
    unit 4 {  
      family inet;  
    }  
  }  
  lo0 {  
    unit 4 {  
      family inet {  
        address 172.16.1.4/32;  
      }  
    }  
    unit 104 {  
      family inet {  
        address 172.16.100.4/32;  
      }  
    }  
  }  
}
```

```
user@R4# show protocols  
  rsvp {  
    interface all {  
      aggregate;  
    }  
  }  
  mpls {  
    interface all;  
    interface ge-1/2/0.10;  
  }  
  bgp {  
    group ibgp {  
      type internal;  
      local-address 172.16.1.4;  
      family inet-vpn {  
        unicast;  
        any;  
      }  
      family inet-mvpn {  
        signaling {  
          damping;  
        }  
      }  
    }  
  }  
}
```

```

    }
  }
  neighbor 172.16.1.2 {
    import dampPolicy;
  }
  neighbor 172.16.1.5;
}
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface lo0.4 {
      passive;
    }
    interface ge-1/2/0.10;
  }
}
}
ldp {
  interface ge-1/2/0.10;
  p2mp;
}
}

```

```

user@R4# show policy-options
policy-statement dampPolicy {
  term term1 {
    from {
      family inet-mvpn;
      nlri-route-type [ 3 4 5 ];
    }
    then accept;
  }
  then {
    damping no-damp;
    accept;
  }
}
policy-statement parent_vpn_routes {
  from protocol bgp;
  then accept;
}
}

```



```
damping no-damp {  
    disable;  
}
```

```
user@R4# show routing-instances  
vpn-1 {  
    instance-type vrf;  
    interface vt-1/2/0.4;  
    interface ge-1/2/1.17;  
    interface lo0.104;  
    route-distinguisher 100:100;  
    vrf-target target:1:1;  
    protocols {  
        ospf {  
            export parent_vpn_routes;  
            area 0.0.0.0 {  
                interface lo0.104 {  
                    passive;  
                }  
                interface ge-1/2/1.17;  
            }  
        }  
        pim {  
            rp {  
                static {  
                    address 172.16.100.2;  
                }  
            }  
            interface ge-1/2/1.17 {  
                mode sparse;  
            }  
        }  
        mvpn;  
    }  
}
```

```
user@R4# show routing-opts  
router-id 172.16.1.4;  
autonomous-system 1001;
```

Verification

IN THIS SECTION

- [Verifying That Route Flap Damping Is Disabled | 832](#)
- [Verifying Route Flap Damping | 833](#)

Confirm that the configuration is working properly.

Verifying That Route Flap Damping Is Disabled

Purpose

Verify the presence of the `no-damp` policy, which disables damping for MVPN route types other than 3, 4, and 5.

Action

From operational mode, enter the `show policy damping` command.

```
user@R4> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
Computed values:
  Merit ceiling: 12110
  Maximum decay: 6193
Damping information for "no-damp":
  Damping disabled
```

Meaning

The output shows that the default damping parameters are in effect and that the `no-damp` policy is also in effect for the specified route types.

Verifying Route Flap Damping

Purpose

Check whether BGP routes have been damped.

Action

From operational mode, enter the `show bgp summary` command.

```

user@R4> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History Damp State  Pending
bgp.l3vpn.0
                6          6          0          0          0          0
bgp.l3vpn.2
                0          0          0          0          0          0
bgp.mvpn.0
                2          2          0          0          0          0
Peer           AS      InPkt   OutPkt   OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
172.16.1.2     1001    3159    3155     0      0    23:43:47 Establ
  bgp.l3vpn.0: 3/3/3/0
  bgp.l3vpn.2: 0/0/0/0
  bgp.mvpn.0:  1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0
172.16.1.5     1001    3157    3154     0      0    23:43:40 Establ
  bgp.l3vpn.0: 3/3/3/0
  bgp.l3vpn.2: 0/0/0/0
  bgp.mvpn.0:  1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0

```

Meaning

The Damp State field shows that zero routes in the `bgp.mvpn.0` routing table have been damped. Further down, the last number in the State field shows that zero routes have been damped for BGP peer `172.16.1.2`.

SEE ALSO

Understanding Damping Parameters

Using Routing Policies to Damp BGP Route Flapping

Example: Configuring BGP Route Flap Damping Parameters

Example: Configuring MBGP Multicast VPN Topology Variations

IN THIS SECTION

- [Requirements | 834](#)
- [Overview and Topology | 834](#)
- [Configuring Full Mesh MBGP MVPNs | 837](#)
- [Configuring Sender-Only and Receiver-Only Sites Using PIM ASM Provider Tunnels | 840](#)
- [Configuring Sender-Only, Receiver-Only, and Sender-Receiver MVPN Sites | 843](#)
- [Configuring Hub-and-Spoke MVPNs | 847](#)

This section describes how to configure multicast virtual private networks (MVPNs) using multiprotocol BGP (MBGP) (next-generation MVPNs).

Requirements

To implement multiprotocol BGP-based multicast VPNs, auto-RP, bootstrap router (BSR) RP, and PIM dense mode you need JUNOS Release 9.2 or later.

To implement multiprotocol BGP-based multicast VPNs, sender-only sites, and receiver-only sites you need JUNOS Release 8.4 or later.

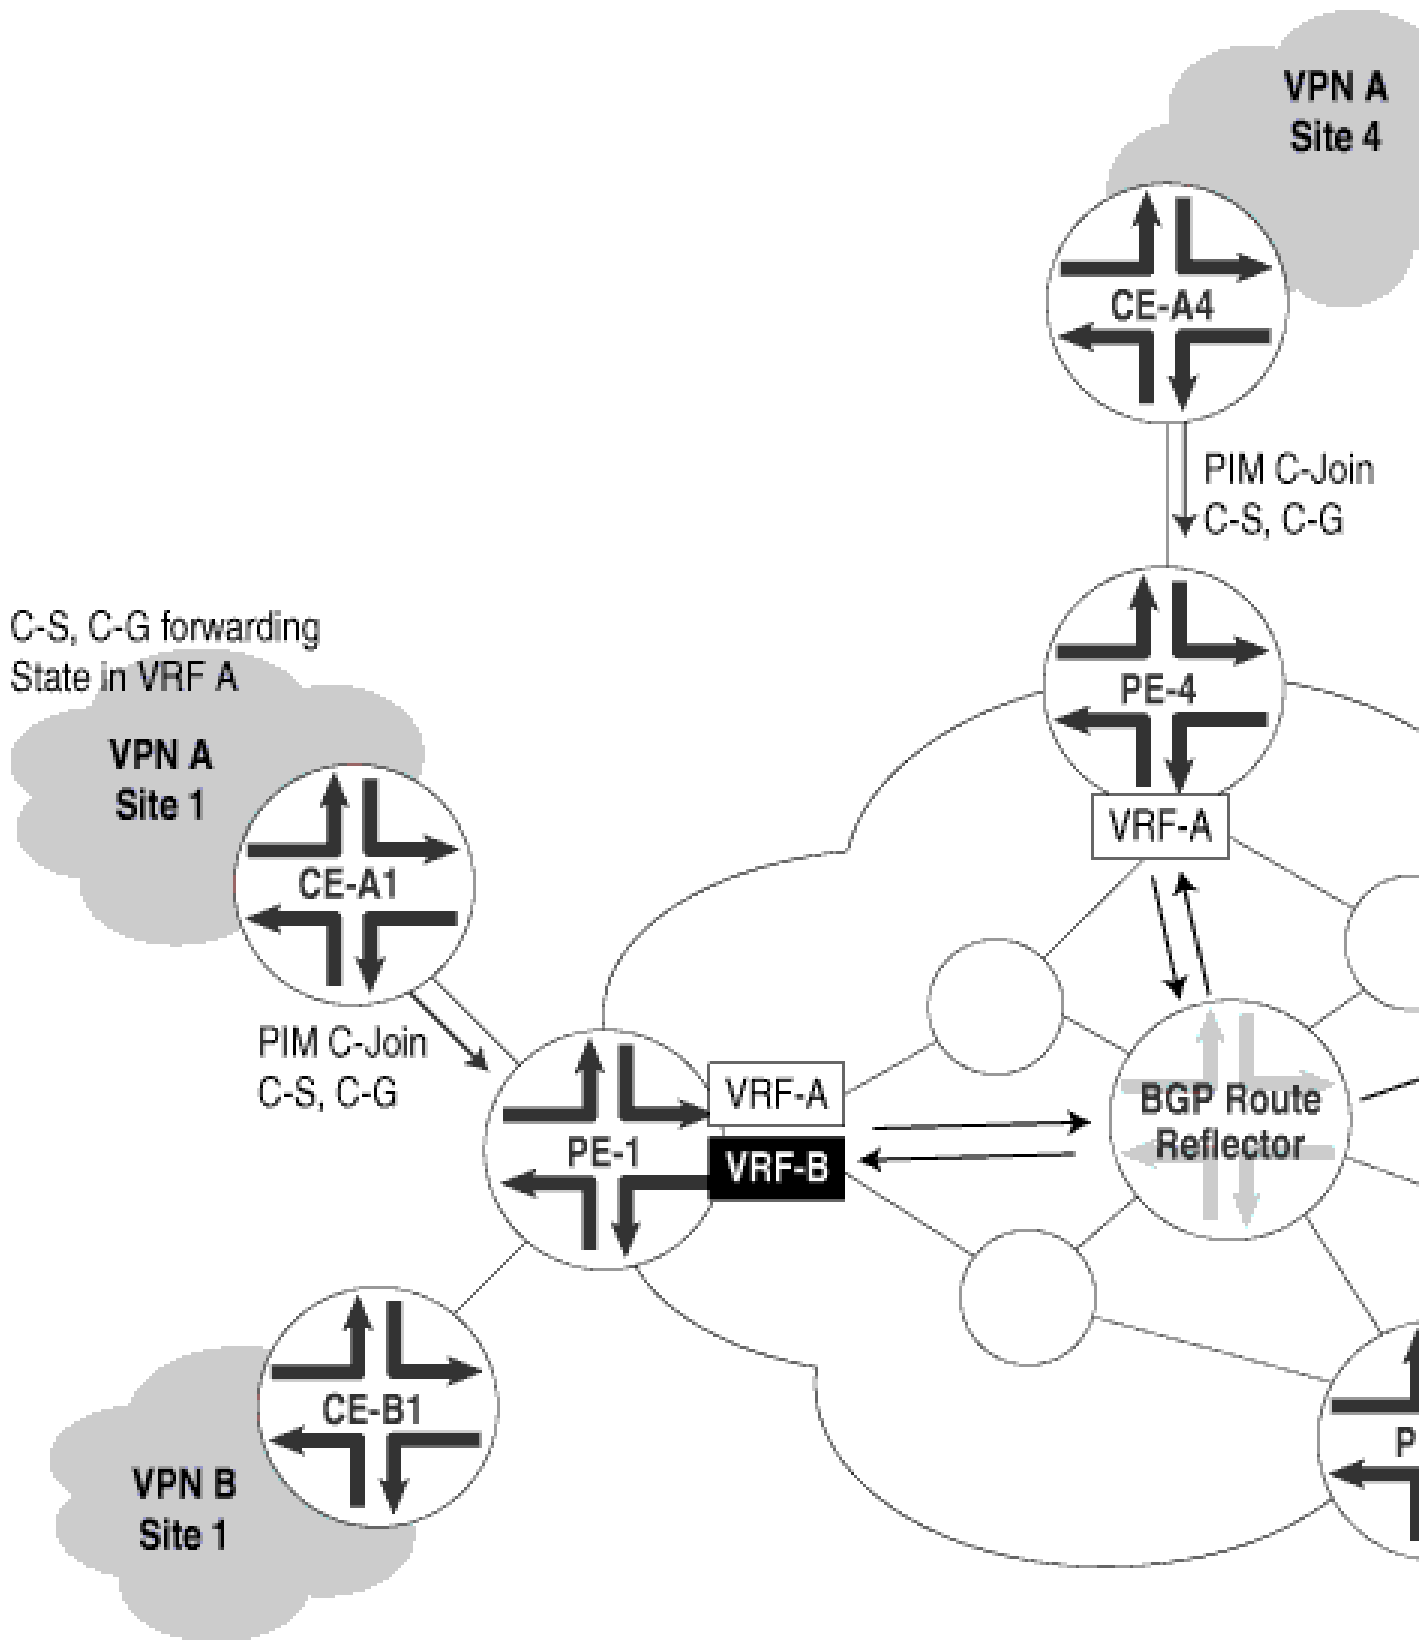
Overview and Topology

You can configure PIM auto-RP, bootstrap router (BSR) RP, PIM dense mode, and mtrace for next generation multicast VPN networks. 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 is the IETF standard for RP establishment. A selected router in a network acts as a BSR, which selects a unique RP for different group ranges. BSR messages are flooded using the data tunnel between PE routers. When you enable PIM dense mode, data packets are forwarded to all interfaces except the incoming interface. Unlike PIM sparse mode, where explicit joins are required for data packets to be transmitted downstream, data packets are flooded to all routers in the routing instance in PIM dense mode.

This section shows you how to configure a MVPN using MBGP. If you have multicast VPNs based on draft-rosen, they will continue to work as before and are not affected by the configuration of MVPNs using MBGP.

The network configuration used for most of the examples in this section is shown in [Figure 111 on page 836](#).

Figure 111: MBGP MVPN Topology Variations Diagram



In the figure, two VPNs, VPN A and VPN B, are serviced by the same provider at several sites, two of which have CE routers for both VPN A and VPN B (site 2 is not shown). The PE routers are shown with VRF tables for the VPN CEs for which they have routing information. It is important to note that no multicast protocols are required between the PE routers on the network. The multicast routing information is carried by MBGP between the PE routers. There may be one or more BGP route reflectors in the network. Both VPNs operate independently and are configured separately.

Both the PE and CE routers run PIM sparse mode and maintain forwarding state information about customer source (C-S) and customer group (C-G) multicast components. CE routers still send a customer's PIM join messages (PIM C-Join) from CE to PE, and from PE to CE, as shown in the figure. But on the provider's backbone network, all multicast information is carried by MBGP. The only addition over and above the unicast VPN configuration normally used is the use of a special provider tunnel (**provider-tunnel**) for carrying PIM sparse mode message content between provider nodes on the network.

There are several scenarios for MVPN configuration using MBGP, depending on whether a customer site has senders (sources) of multicast traffic, has receivers of multicast traffic, or a mixture of senders and receivers. MVPNs can be:

- A full mesh (each MVPN site has both senders and receivers)
- A mixture of sender-only and receiver-only sites
- A mixture of sender-only, receiver-only, and sender-receiver sites
- A hub and spoke (two interfaces between hub PE and hub CE, and all spokes are sender-receiver sites)

Each type of MVPN differs more in the configuration VPN statements than the provider tunnel configuration. For information about configuring VPNs, see the [Junos OS VPNs Library for Routing Devices](#).

Configuring Full Mesh MBGP MVPNs

IN THIS SECTION

- [Configuration Steps | 838](#)

This example describes how to configure a full mesh MBGP MVPN:

Configuration Steps

Step-by-Step Procedure

In this example, PE-1 connects to VPN A and VPN B at site 1, PE-4 connects to VPN A at site 4, and PE-2 connects to VPN B at site 3. To configure a full mesh MVPN for VPN A and VPN B, perform the following steps:

1. Configure PE-1 (both VPN A and VPN B at site 1):

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface so-6/0/0.0;
    interface so-6/0/1.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.1;
      }
    }
    protocols {
      mvpn;
    }
    route-distinguisher 65535:0;
    vrf-target target:1:1;
  }
  VPN-B {
    instance-type vrf;
    interface ge-0/3/0.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.2;
      }
    }
    protocols {
      mvpn;
    }
    route-distinguisher 65535:1;
    vrf-target target:1:2;
  }
}
```


2. Configure PE-4 (VPN A at site 4):

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface so-1/0/0.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.1;
      }
    }
    protocols {
      mvpn;
    }
    route-distinguisher 65535:4;
    vrf-target target:1:1;
  }
}
```

3. Configure PE-2 (VPN B at site 3):

```
[edit]
routing-instances {
  VPN-B {
    instance-type vrf;
    interface ge-1/3/0.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.2;
      }
    }
    protocols {
      mvpn;
    }
    route-distinguisher 65535:3;
    vrf-target target:1:2;
  }
}
```

Configuring Sender-Only and Receiver-Only Sites Using PIM ASM Provider Tunnels

IN THIS SECTION

- [Configuration Steps | 840](#)

This example describes how to configure an MBGP MVPN with a mixture of sender-only and receiver-only sites using PIM-ASM provider tunnels.

Configuration Steps

Step-by-Step Procedure

In this example, PE-1 connects to VPN A (sender-only) and VPN B (receiver-only) at site 1, PE-4 connects to VPN A (receiver-only) at site 4, and PE-2 connects to VPN A (receiver-only) and VPN B (sender-only) at site 3.

To configure an MVPN for a mixture of sender-only and receiver-only sites on VPN A and VPN B, perform the following steps:

1. Configure PE-1 (VPN A sender-only and VPN B receiver-only at site 1):

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface so-6/0/0.0;
    interface so-6/0/1.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.1;
      }
    }
  }
  protocols {
    mvpn {
      sender-site;
      route-target {
        export-target unicast;
        import-target target target:1:4;
      }
    }
  }
}
```

```

    }

    }
}
route-distinguisher 65535:0;
vrf-target target:1:1;
routing-options {
    auto-export;
}
}
VPN-B {
    instance-type vrf;
    interface ge-0/3/0.0;
    provider-tunnel {
        pim-asm {
            group-address 224.1.1.2;
        }
    }
    protocols {
        mvpn {
            receiver-site;
            route-target {
                export-target target target:1:5;
                import-target unicast;
            }
        }
    }
    route-distinguisher 65535:1;
    vrf-target target:1:2;
    routing-options {
        auto-export;
    }
}
}

```

2. Configure PE-4 (VPN A receiver-only at site 4):

```

[edit]
routing-instances {
    VPN-A {
        instance-type vrf;
        interface so-1/0/0.0;
        provider-tunnel {

```

```

    pim-asm {
        group-address 224.1.1.1;
    }
}
protocols {
    mvpn {
        receiver-site;
        route-target {
            export-target target target:1:4;
            import-target unicast;
        }
    }
}
route-distinguisher 65535:2;
vrf-target target:1:1;
routing-options {
    auto-export;
}
}

```

3. Configure PE-2 (VPN A receiver-only and VPN B sender-only at site 3):

```

[edit]
routing-instances {
    VPN-A {
        instance-type vrf;
        interface so-2/0/1.0;
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
        }
    }
    protocols {
        mvpn {
            receiver-site;
            route-target {
                export-target target target:1:4;
                import-target unicast;
            }
        }
    }
}
}

```

```

route-distinguisher 65535:3;
vrf-target target:1:1;
routing-options {
    auto-export;
}
}
VPN-B {
    instance-type vrf;
    interface ge-1/3/0.0;
    provider-tunnel {
        pim-asm {
            group-address 224.1.1.2;
        }
    }
    protocols {
        mvpn {
            sender-site;
            route-target {
                export-target unicast
                import-target target target:1:5;
            }
        }
    }
    route-distinguisher 65535:4;
    vrf-target target:1:2;
    routing-options {
        auto-export;
    }
}

```

Configuring Sender-Only, Receiver-Only, and Sender-Receiver MVPN Sites

IN THIS SECTION

- [Configuration Steps | 844](#)

This example describes how to configure an MBGP MVPN with a mixture of sender-only, receiver-only, and sender-receiver sites.

Configuration Steps

Step-by-Step Procedure

In this example, PE-1 connects to VPN A (sender-receiver) and VPN B (receiver-only) at site 1, PE-4 connects to VPN A (receiver-only) at site 4, and PE-2 connects to VPN A (sender-only) and VPN B (sender-only) at site 3. To configure an MVPN for a mixture of sender-only, receiver-only, and sender-receiver sites for VPN A and VPN B, perform the following steps:

1. Configure PE-1 (VPN A sender-receiver and VPN B receiver-only at site 1):

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface so-6/0/0.0;
    interface so-6/0/1.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.1;
      }
    }
    protocols {
      mvpn {
        route-target {
          export-target unicast target target:1:4;
          import-target unicast target target:1:4 receiver;
        }
      }
    }
    route-distinguisher 65535:0;
    vrf-target target:1:1;
    routing-options {
      auto-export;
    }
  }
  VPN-B {
    instance-type vrf;
    interface ge-0/3/0.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.2;
      }
    }
  }
}
```

```

    }
  }
  protocols {
    mvpn {
      receiver-site;
      route-target {
        export-target target target:1:5;
        import-target unicast;
      }
    }
  }
  route-distinguisher 65535:1;
  vrf-target target:1:2;
  routing-options {
    auto-export;
  }
}

```

2. Configure PE-4 (VPN A receiver-only at site 4):

```

[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface so-1/0/0.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.1;
      }
    }
    protocols {
      mvpn {
        receiver-site;
        route-target {
          export-target target target:1:4;
          import-target unicast;
        }
      }
    }
  }
  route-distinguisher 65535:2;
  vrf-target target:1:1;
}

```

```

    routing-options {
      auto-export;
    }
  }
}

```

3. Configure PE-2 (VPN-A sender-only and VPN-B sender-only at site 3):

```

[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface so-2/0/1.0;
    provider-tunnel {
      pim-asm {
        group-address 224.1.1.1;
      }
    }
    protocols {
      mvpn {
        receiver-site;
        route-target {
          export-target target target:1:4;
          import-target unicast;
        }
      }
    }
  }
  route-distinguisher 65535:3;
  vrf-target target:1:1;
  routing-options {
    auto-export;
  }
}
VPN-B {
  instance-type vrf;
  interface ge-1/3/0.0;
  provider-tunnel {
    pim-asm {
      group-address 224.1.1.2;
    }
  }
  protocols {

```



```

    mvpn {
        sender-site;
        route-target {
            export-target unicast;
            import-target target target:1:5;
        }
    }
}
route-distinguisher 65535:4;
vrf-target target:1:2;
routing-options {
    auto-export;
}
}

```

Configuring Hub-and-Spoke MVPNs

IN THIS SECTION

- [Configuration Steps | 847](#)

This example describes how to configure an MBGP MVPN in a hub and spoke topology.

Configuration Steps

Step-by-Step Procedure

In this example, which only configures VPN A, PE-1 connects to VPN A (spoke site) at site 1, PE-4 connects to VPN A (hub site) at site 4, and PE-2 connects to VPN A (spoke site) at site 3. Current support is limited to the case where there are two interfaces between the hub site CE and PE. To configure a hub-and-spoke MVPN for VPN A, perform the following steps:

1. Configure PE-1 for VPN A (spoke site):

```

[edit]
routing-instances {
    VPN-A {
        instance-type vrf;
    }
}

```

```

interface so-6/0/0.0;
interface so-6/0/1.0;
provider-tunnel {
    rsvp-te {
        label-switched-path-template {
            default-template;
        }
    }
}
protocols {
    mvpn {
        route-target {
            export-target unicast;
            import-target unicast target target:1:4;
        }
    }
}
route-distinguisher 65535:0;
vrf-target {
    import target:1:1;
    export target:1:3;
}
routing-options {
    auto-export;
}
}

```

2. Configure PE-4 for VPN A (hub site):

```

[edit]
routing-instances {
    VPN-A-spoke-to-hub {
        instance-type vrf;
        interface so-1/0/0.0; #receives data and joins from the CE
        protocols {
            mvpn {
                receiver-site;
                route-target {
                    export-target target target:1:4;
                    import-target unicast;
                }
            }
        }
    }
}

```

```

    ospf {
        export redistribute-vpn; #redistributes VPN routes to CE
        area 0.0.0.0 {
            interface so-1/0/0;
        }
    }
}
route-distinguisher 65535:2;
vrf-target {
    import target:1:3;
}
routing-options {
    auto-export;
}
}
VPN-A-hub-to-spoke {
    instance-type vrf;
    interface so-2/0/0.0; #receives data and joins from the CE
    provider-tunnel {
        rsvp-te {
            label-switched-path-template {
                default-template;
            }
        }
    }
}
protocols {
    mvpn {
        sender-site;
        route-target {
            import-target target target:1:3;
            export-target unicast;
        }
    }
    ospf {
        export redistribute-vpn; #redistributes VPN routes to CE
        area 0.0.0.0 {
            interface so-2/0/0;
        }
    }
}
route-distinguisher 65535:2;
vrf-target {
    import target:1:1;
}

```

```

    }
    routing-options {
        auto-export;
    }
}

```

3. Configure PE-2 for VPN A (spoke site):

```

[edit]
routing-instances {
    VPN-A {
        instance-type vrf;
        interface so-2/0/1.0;
        provider-tunnel {
            rsvp-te {
                label-switched-path-template {
                    default-template;
                }
            }
        }
        protocols {
            mvpn {
                route-target {
                    import-target target target:1:4;
                    export-target unicast;
                }
            }
        }
        route-distinguisher 65535:3;
        vrf-target {
            import target:1:1;
            export target:1:3;
        }
        routing-options {
            auto-export;
        }
    }
}

```

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.

Before you begin:

- Configure the router interfaces. See [Interfaces Fundamentals](#).
- 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.

The state maintained by MVPN includes MVPN routes, cmcast, provider-tunnel, and forwarding information. BGP MVPN NSR synchronizes this MVPN state between the primary 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

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 designated router 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 designated router.

```
[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 designated router.

```
[edit]
user@host# set routing-options forwarding-table export load-balance
```

6. Configure nonstop active routing on the designated router.

```
[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
```

SEE ALSO

[Configuring Basic PIM Settings](#)

[Understanding Nonstop Active Routing for PIM](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.1X49-D50	Starting in Junos OS Release 15.1X49-D50 and Junos OS Release 17.3R1, the <code>vrf-table-label</code> statement allows mapping of the inner label to a specific Virtual Routing and Forwarding (VRF). This mapping allows examination of the encapsulated IP header at an egress VPN router. For SRX Series Firewalls, the <code>vrf-table-label</code> statement is currently supported only on physical interfaces. As a workaround, deactivate <code>vrf-table-label</code> or use physical interfaces.

RELATED DOCUMENTATION

[Example: Configuring MBGP MVPN Extranets | 858](#)

[Multiprotocol BGP MVPNs Overview | 741](#)

BGP-MVPN Inter-AS Option B Overview

This topic provides an overview of Junos support for Inter-Autonomous System (AS) Option B, which is achieved by extending Border Gateway Protocol Multicast Virtual Private Network (BGP-MVPN) to support Inter-AS scenarios using segmented provider tunnels (p-tunnels). Junos OS also support Option A and Option C unicast with non-segmented p-tunnels, support for which was introduced in Junos OS 12.1. See the links below for more information on these options.

Inter-AS support for multicast traffic is required when an L3VPN results in two or more ASes that are using BGP-MVPN. The ASes may be administered by the same authority or by different authorities. When using BGP-MVPN Inter-AS Option B with segmented p-tunnels, the p-tunnel segmentation is performed at the Autonomous System Border Router (ASBRs). The ASBRs also perform BGP-MVPN signaling and form the data plane.

Setting up Inter-AS Option B with segmented p-tunnels can be complex, but the configuration does provide the following advantages:

- Independence. Different administrative authorities can choose whether or not to allow topology discovery of their AS by the other ASes. That is, each AS can be separately controlled by a different independent authority.
- Heterogeneity. Different p-tunnel technologies can be used within a given AS (as might be the case when working with heterogeneous networks that now must be combined).

- Scale. Inter-AS Option B with segmented p-tunnels avoids the potential for ASBR bottleneck that can happen when Intra-AS p-tunnels are set up across ASes using non-segmented p-tunnels. (Unicast branch LSPs with inclusive p-tunnels can all have to transit through the ASBRs. In this case, for IR, the pinch point becomes data-plane scale. For RSVP-TE it becomes P2MP control-plane scale, due to the high number of RSVP refresh messages passing through the ASBRs).

The supported Junos implementation of Option B uses RSVP-TE p-tunnels for all segments, and MVPN Inter-AS signaling procedures. Multicast traffic is forwarded across AS boundaries over a single-hop labeled LSP. Inter-AS p-tunnels have two segments: an ASBR-ASBR segment, called *Inter-AS* segment and the ASBR-PE segment called *Intra-AS* segment. (Static RSVP-TE, IR, PIM-ASM, and PIM-SSM p-tunnels are not supported.)

MVPN Intra-AS AD routes are not propagated across the AS boundary. The Intra-AS inclusive p-tunnels advertised in Type-1 routes are terminated at the ASBRs within each AS. Route learning for both unicast and multicast traffic occurs only through Option B.

The ASBR originates an Inter-AS AD (Type-2) route into eBGP, which may include tunnel attributes for an Inter-AS p-tunnel (called an Inter-AS, or ASBR-ASBR p-tunnel segment). The Type-2 route contains the ASBR's route distinguisher (RD), which is unique per VPN and per ASBR, and its AS number. The tunnel is set up between two directly connected ASBRs in neighboring ASes, and it is always a single-hop point-to-point (P2P) LSP.

An ASBR in the originating AS forwards all multicast traffic received over the inclusive p-tunnel into the Inter-AS p-tunnel. An ASBR in the adjacent AS propagates the received Inter-AS route into its own AS over iBGP, but only after rewriting the Provider Multicast Service Interface (PMSI) tunnel attributes and modifying the next-hop of the Multiprotocol Reachable (MP_REACH_NRLI) attribute with a reachable address of the ASBR (next-hop self rewrite). When an ASBR propagates the Type-2 route over iBGP, it can choose any p-tunnel type supported within its AS, although the supported Junos implementation of Option B uses RSVP-TE p-tunnels only for all segments.

At the ASBRs, traffic received over the upstream p-tunnel segment is forwarded over the downstream p-tunnel segment. This process is repeated at each AS boundary. The resulting Inter-AS p-tunnel is comprised of alternating Inter-AS and Intra-AS p-tunnel segments (thus the name, "segmented p-tunnel").

Option B with segmented p-tunnels is not without drawbacks.:

- The ASBRs distribute both VPN routes and routes in the master instance. They may thus become a bottleneck.
- With a large number of VPNs, the ASBR can run out of labels because each unicast VPN route requires one.
- Per VPN packet flow accounting cannot be performed at the ASBR.

- Unless route-targets are rewritten at the AS boundaries, the different service providers must agree on VPN route-targets (this is that same as for option-C)
- The ASBRs must be capable of MVPN signaling and support Inter-AS MVPN procedures.

RELATED DOCUMENTATION

| [inter-as](#)

ACX Support for BGP MVPN

ACX7100-32C and ACX7100-48L routers running Junos Evolved release 22.1R1 support BGP/MPLS MVPN (also known as “next generation,” or “NG,” MVPN) running on multipoint LDP (MLDP P2MP) provider tunnels, where BGP MVPN is the intra-AS and PIM-SM/SSM is the data plane.

The following considerations apply:

- BGP intra-AS auto-discovery occurs at the level of PE, MVPN and the binding of a MVPN to an Inclusive Tree.
 - For a PE in a MVPN *sender* site set to discover PEs in the MVPN *receiver* sites set, the AD route is imported by a PE attached to a *sender* site into the VRF only if the import RT matches the AD route RT. A PE in the *receiver* sites set exports the receiver sites RT.
 - For a PE in a MVPN *receiver* sites set to discover PEs in the MVPN *sender* sites set, the AD route is imported by a PE attached to a *receiver* site into the VRF only if the import RT matches the AD route RT. A PE in the *sender* sites set exports the sender sites RT.
 - Constrained set up of *Inclusive* P-Tunnels and the binding of a MVPN to such tunnels to construct an I-PMSI. A PE in the sender site set of a particular MVPN advertises an intra-AS AD route that carries the P-Tunnel identifier in a new attribute. The PE sends packets received from the sender sites of this MVPN on the P-Tunnel. A PE in the receiver site set of this MVPN associates this P-tunnel with the VRF. The PE in the receiver site forwards packets received on the P-tunnel using the VRF.
- MLDP Inclusive Intra-AS P-Tunnels. MLDP P2MP is supported as the data plane for Inclusive P-Tunnels. Aggregate trees are not supported, thus a separate P-tunnel is setup for a given MVPN.
 - A provider must configure MLDP P2MP as the P-Tunnel technology on each PE in the MVPN that belongs to the sender site set. By default a PE belongs to both sender site set and receiver site set. Configuration can be used to change the default behavior.
 - A provider must configure the same P-Tunnel MLDP P2MP technology on all PEs in the MVPN.

- A provider must also configure MLDP P2MP on all P routers and MUST configure in the PEs.
- The goal is to provide a label-based switching in the core for the multicast traffic from the CE site to the other CE site.
- MLDP MP2MP Inclusive Intra-AS P-Tunnels are not supported in release 22.1R1.
- C-multicast Routing Information Exchange using BGP provides the intra-AS C-multicast routing exchange using BGP (as covered in BGP-MVPN) and which provides the following:
 - Constrained exchange of C-multicast routes using BGP is done by:
 - Constraining distribution of unicast VPN-IPv4 routes. A PE connected to a receiver site cannot join a C-S unless it has a unicast route to C-S, which is handled by the unicast (VPN-IPv4) routing.
 - Using C-multicast routes to pass MVPN C-multicast routing information from PEs connected to sites in receiver sites set to PEs connected to sites in sender sites set.
 - Using Route Import Extended Community to control import of C-multicast routes into a particular VRF [BGP-MVPN]. The Route Import Extended Community is distinct per VRF. It contains the PE's IP address and number (which are local to the PE). C-multicast route is accepted into a given VRF only if the route target carried by the C-multicast route is equal to the route import associated with the VRF.
 - Supporting MVPN customers who use PIM-SM in the SSM mode. This requires only one type of C-multicast routes, the Source Tree Join.
 - Supporting MVPN customers who use PIM-SM in the SM mode. This is provided by using only one type of C-multicast route, the Source Tree Join, which can be enabled because:
 - Each MVPN uses the existing IP multicast mechanisms to enable C-RPs in the PEs.
 - The PEs maintain information about active sources (S,G) for a given MVPN.
 - BGP AD routes are used to distribute information about active sources (S,G) within a given MVPN among all PEs with that MVPN.
 - BGP Source Tree Join C-multicast routes are used to inform PEs connected to active sources that there are receivers for (S,G) connected to some other PEs.

RELATED DOCUMENTATION

<https://datatracker.ietf.org/doc/html/rfc4364>

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/concept/ng-mvpn-control-plane.html>

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/concept/ng-mvpn-topology.html>

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-multicast.html#id-example-configuring-mbgp-multicast-vpns>

Example: Configuring MBGP MVPN Extranets

IN THIS SECTION

- [Understanding MBGP Multicast VPN Extranets | 858](#)
- [MBGP Multicast VPN Extranets Configuration Guidelines | 859](#)
- [Example: Configuring MBGP Multicast VPN Extranets | 860](#)

Understanding MBGP Multicast VPN Extranets

IN THIS SECTION

- [MBGP Multicast VPN Extranets Application | 859](#)

A multicast VPN (MVPN) extranet enables service providers to forward IP multicast traffic originating in one VPN routing and forwarding (VRF) instance to receivers in a different VRF instance. This capability is also known as *overlapping* MVPNs.

The MVPN extranet feature supports the following traffic flows:

- A receiver in one VRF can receive multicast traffic from a source connected to a different router in a different VRF.
- A receiver in one VRF can receive multicast traffic from a source connected to the same router in a different VRF.
- A receiver in one VRF can receive multicast traffic from a source connected to a different router in the same VRF.

- A receiver in one VRF can be prevented from receiving multicast traffic from a specific source in a different VRF.

MBGP Multicast VPN Extranets Application

An MVPN extranet is useful in the following applications.

Mergers and Data Sharing

An MVPN extranet is useful when there are business partnerships between different enterprise VPN customers that require them to be able to communicate with one another. For example, a wholesale company might want to broadcast inventory to its contractors and resellers. An MVPN extranet is also useful when companies merge and one set of VPN sites needs to receive content from another VPN. The enterprises involved in the merger are different VPN customers from the service provider point of view. The MVPN extranet makes the connectivity possible.

Video Distribution

Another use for MVPN extranets is video multicast distribution from a video headend to receiving sites. Sites within a given multicast VPN might be in different organizations. The receivers can subscribe to content from a specific content provider.

The PE routers on the MVPN provider network learn about the sources and receivers using MVPN mechanisms. These PE routers can use selective trees as the multicast distribution mechanism in the backbone. The network carries traffic belonging only to a specified set of one or more multicast groups, from one or more multicast VPNs. As a result, this model facilitates the distribution of content from multiple providers on a selective basis if desired.

Financial Services

A third use for MVPN extranets is enterprise and financial services infrastructures. The delivery of financial data, such as financial market updates, stock ticker values, and financial TV channels, is an example of an application that must deliver the same data stream to hundreds and potentially thousands of end users. The content distribution mechanisms largely rely on multicast within the financial provider network. In this case, there could also be an extensive multicast topology within brokerage firms and banks networks to enable further distribution of content and for trading applications. Financial service providers require traffic separation between customers accessing the content, and MVPN extranets provide this separation.

MBGP Multicast VPN Extranets Configuration Guidelines

When configuring MVPN extranets, keep the following in mind:

- If there is more than one VRF routing instance on a provider edge (PE) router that has receivers interested in receiving multicast traffic from the same source, virtual tunnel (VT) interfaces must be configured on all instances.

- For auto-RP operation, the mapping agent must be configured on at least two PEs in the extranet network.
- For asymmetrically configured extranets using auto-RP, when one VRF instance is the only instance that imports routes from all other extranet instances, the mapping agent must be configured in the VRF that can receive all RP discovery messages from all VRF instances, and mapping-agent election should be disabled.
- For bootstrap router (BSR) operation, the candidate and elected BSRs can be on PE, CE, or C routers. The PE router that connects the BSR to the MVPN extranets must have configured provider tunnels or other physical interfaces configured in the routing instance. The only case not supported is when the BSR is on a CE or C router connected to a PE routing instance that is part of an extranet but does not have configured provider tunnels and does not have any other interfaces besides the one connecting to the CE router.
- RSVP-TE point-to-multipoint LSPs must be used for the provider tunnels.
- PIM dense mode is not supported in the MVPN extranets VRF instances.

Example: Configuring MBGP Multicast VPN Extranets

IN THIS SECTION

- [Requirements | 860](#)
- [Overview and Topology | 861](#)
- [Configuration | 862](#)

This example provides a step-by-step procedure to configure multicast VPN extranets using static rendezvous points. It is organized in the following sections:

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later
- Six T Series, or MX Series Juniper routers
- One adaptive services PIC or MultiServices PIC in each of the T Series routers acting as PE routers
- One host system capable of sending multicast traffic and supporting the Internet Group Management Protocol (IGMP)

- Three host systems capable of receiving multicast traffic and supporting IGMP

Overview and Topology

IN THIS SECTION

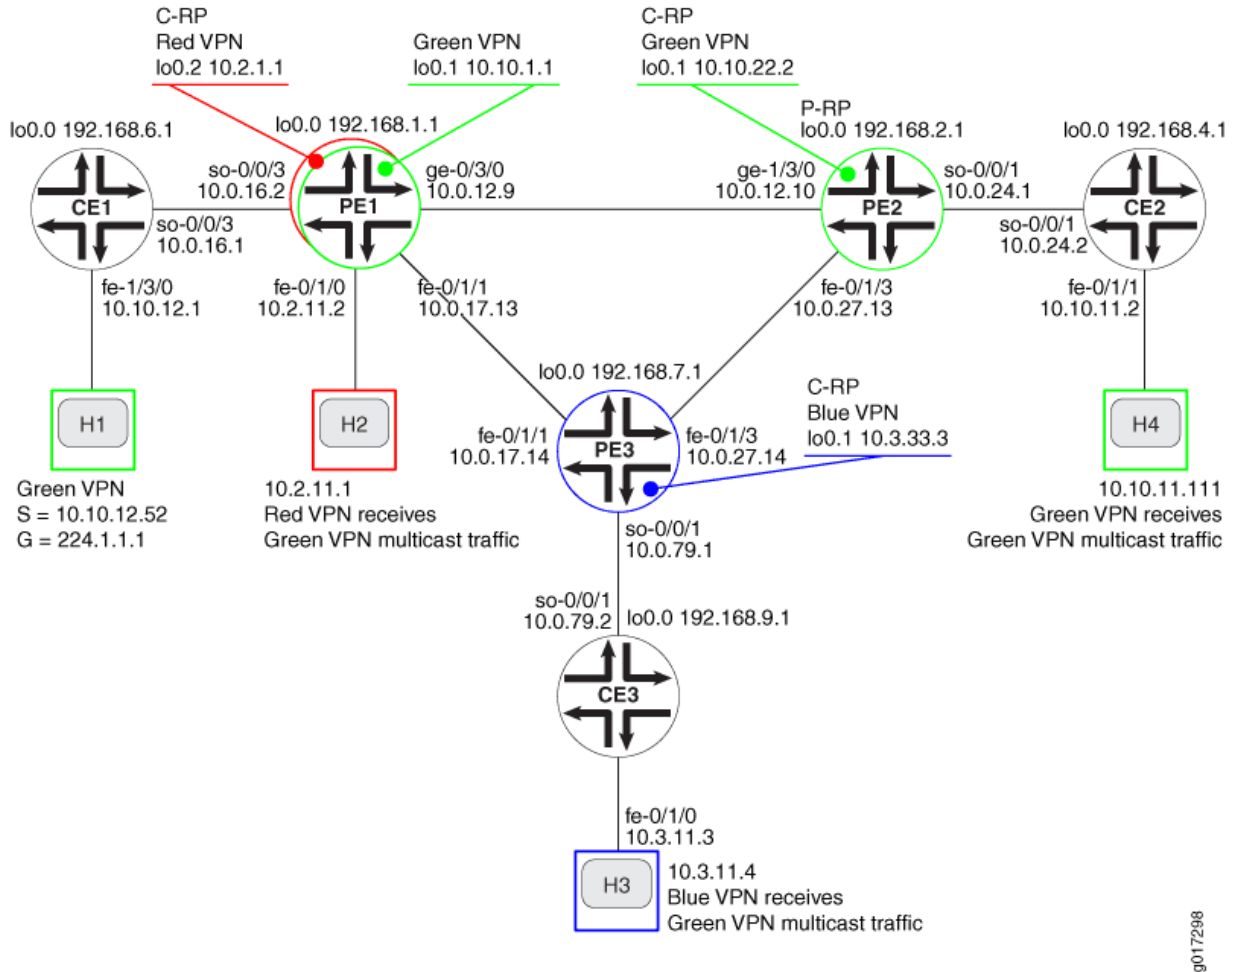
- [Topology | 862](#)

In the network topology shown in [Figure 112 on page 862](#):

- Host H1 is the source for group 244.1.1.1 in the green VPN.
- The multicast traffic originating at source H1 can be received by host H4 connected to router CE2 in the green VPN.
- The multicast traffic originating at source H1 can be received by host H3 connected to router CE3 in the blue VPN.
- The multicast traffic originating at source H1 can be received by host H2 directly connected to router PE1 in the red VPN.
- Any host can be a sender site or receiver site.

Topology

Figure 112: MVPN Extranets Topology Diagram



Configuration

IN THIS SECTION

- [Configuring Interfaces | 863](#)
- [Configuring an IGP in the Core | 866](#)
- [Configuring BGP in the Core | 867](#)
- [Configuring LDP | 869](#)
- [Configuring RSVP | 871](#)

- [Configuring MPLS | 872](#)
- [Configuring the VRF Routing Instances | 873](#)
- [Configuring MVPN Extranet Policy | 876](#)
- [Configuring CE-PE BGP | 881](#)
- [Configuring PIM on the PE Routers | 883](#)
- [Configuring PIM on the CE Routers | 885](#)
- [Configuring the Rendezvous Points | 886](#)
- [Testing MVPN Extranets | 889](#)
- [Results | 892](#)



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:

- CE1 identifies the customer edge 1 (CE1) router
- PE1 identifies the provider edge 1 (PE1) router
- CE2 identifies the customer edge 2 (CE2) router
- PE2 identifies the provider edge 2 (PE2) router
- CE3 identifies the customer edge 3 (CE3) router
- PE3 identifies the provider edge 3 (PE3) router

Configuring multicast VPN extranets, involves the following tasks:

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 [Junos OS CLI User Guide](#).

1. On each router, configure an IP address on the loopback logical interface 0 (lo0.0).

```

user@CE1# set interfaces lo0 unit 0 family inet address 192.168.6.1/32 primary
user@PE1# set interfaces lo0 unit 0 family inet address 192.168.1.1/32 primary
user@PE2# set interfaces lo0 unit 0 family inet address 192.168.2.1/32 primary
user@CE2# set interfaces lo0 unit 0 family inet address 192.168.4.1/32 primary
user@PE3# set interfaces lo0 unit 0 family inet address 192.168.7.1/32 primary
user@CE3# set interfaces lo0 unit 0 family inet address 192.168.9.1/32 primary

```

Use the `show interfaces terse` command to verify that the correct IP address is configured on the loopback interface.

2. On the PE and CE routers, configure the IP address and protocol family on the Fast Ethernet and Gigabit Ethernet interfaces. Specify the `inet` address family type.

```

user@CE1# set interfaces fe-1/3/0 unit 0 family inet address 10.10.12.1/24
user@PE1# set interfaces fe-0/1/0 unit 0 description "to H2"
user@PE1# set interfaces fe-0/1/0 unit 0 family inet address 10.2.11.2/30
user@PE1# set interfaces fe-0/1/1 unit 0 description "to PE3 fe-0/1/1.0"
user@PE1# set interfaces fe-0/1/1 unit 0 family inet address 10.0.17.13/30
user@PE1# set interfaces ge-0/3/0 unit 0 family inet address 10.0.12.9/30
user@PE2# set interfaces fe-0/1/3 unit 0 description "to PE3 fe-0/1/3.0"
user@PE2# set interfaces fe-0/1/3 unit 0 family inet address 10.0.27.13/30
user@PE2# set interfaces ge-1/3/0 unit 0 description "to PE1 ge-0/3/0.0"
user@PE2# set interfaces ge-1/3/0 unit 0 family inet address 10.0.12.10/30
user@CE2# set interfaces fe-0/1/1 unit 0 description "to H4"
user@CE2# set interfaces fe-0/1/1 unit 0 family inet address 10.10.11.2/24
user@PE3# set interfaces fe-0/1/1 unit 0 description "to PE1 fe-0/1/1.0"
user@PE3# set interfaces fe-0/1/1 unit 0 family inet address 10.0.17.14/30
user@PE3# set interfaces fe-0/1/3 unit 0 description "to PE2 fe-0/1/3.0"
user@PE3# set interfaces fe-0/1/3 unit 0 family inet address 10.0.27.14/30
user@CE3# set interfaces fe-0/1/0 unit 0 description "to H3"
user@CE3# set interfaces fe-0/1/0 unit 0 family inet address 10.3.11.3/24

```

Use the `show interfaces terse` command to verify that the correct IP address and address family type are configured on the interfaces.

3. On the PE and CE routers, configure the SONET interfaces. Specify the `inet` address family type, and local IP address.

```
user@CE1# set interfaces so-0/0/3 unit 0 description "to PE1 so-0/0/3.0;"
user@CE1# set interfaces so-0/0/3 unit 0 family inet address 10.0.16.1/30
user@PE1# set interfaces so-0/0/3 unit 0 description "to CE1 so-0/0/3.0"
user@PE1# set interfaces so-0/0/3 unit 0 family inet address 10.0.16.2/30
user@PE2# set interfaces so-0/0/1 unit 0 description "to CE2 so-0/0/1:0.0"
user@PE2# set interfaces so-0/0/1 unit 0 family inet address 10.0.24.1/30
user@CE2# set interfaces so-0/0/1 unit 0 description "to PE2 so-0/0/1"
user@CE2# set interfaces so-0/0/1 unit 0 family inet address 10.0.24.2/30
user@PE3# set interfaces so-0/0/1 unit 0 description "to CE3 so-0/0/1.0"
user@PE3# set interfaces so-0/0/1 unit 0 family inet address 10.0.79.1/30
user@CE3# set interfaces so-0/0/1 unit 0 description "to PE3 so-0/0/1"
user@CE3# set interfaces so-0/0/1 unit 0 family inet address 10.0.79.2/30
```

Use the `show configuration interfaces` command to verify that the correct IP address and address family type are configured on the interfaces.

4. On each router, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

5. Use the `ping` command to verify unicast connectivity between each:
 - CE router and the attached host
 - CE router and the directly attached interface on the PE router
 - PE router and the directly attached interfaces on the other PE routers

Configuring an IGP in the Core

Step-by-Step Procedure

On the PE routers, configure an interior gateway protocol such as OSPF or IS-IS. This example shows how to configure OSPF.

1. Specify the `lo0.0` and SONET core-facing logical interfaces.

```
user@PE1# set protocols ospf area 0.0.0.0 interface ge-0/3/0.0 metric 100
user@PE1# set protocols ospf area 0.0.0.0 interface fe-0/1/1.0 metric 100
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.0 passive
user@PE1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
user@PE2# set protocols ospf area 0.0.0.0 interface fe-0/1/3.0 metric 100
user@PE2# set protocols ospf area 0.0.0.0 interface ge-1/3/0.0 metric 100
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.0 passive
user@PE2# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
user@PE3# set protocols ospf area 0.0.0.0 interface lo0.0 passive
user@PE3# set protocols ospf area 0.0.0.0 interface fe-0/1/3.0 metric 100
user@PE3# set protocols ospf area 0.0.0.0 interface fe-0/1/1.0 metric 100
user@PE3# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

2. On the PE routers, configure a router ID.

```
user@PE1# set routing-options router-id 192.168.1.1
user@PE2# set routing-options router-id 192.168.2.1
user@PE3# set routing-options router-id 192.168.7.1
```

Use the `show ospf overview` and `show configuration protocols ospf` commands to verify that the correct interfaces have been configured for the OSPF protocol.

3. On the PE routers, configure OSPF traffic engineering support. Enabling traffic engineering extensions supports the Constrained Shortest Path First algorithm, which is needed to support Resource Reservation Protocol - Traffic Engineering (RSVP-TE) point-to-multipoint label-switched paths (LSPs). If you are configuring IS-IS, traffic engineering is supported without any additional configuration.

```
user@PE1# set protocols ospf traffic-engineering
user@PE2# set protocols ospf traffic-engineering
user@PE3# set protocols ospf traffic-engineering
```

Use the `show ospf overview` and `show configuration protocols ospf` commands to verify that traffic engineering support is enabled for the OSPF protocol.

4. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

5. On the PE routers, verify that the OSPF neighbors form adjacencies.

```
user@PE1> show ospf neighbors
```

Address	Interface	State	ID	Pri	Dead
10.0.17.14	fe-0/1/1.0	Full	192.168.7.1	128	32
10.0.12.10	ge-0/3/0.0	Full	192.168.2.1	128	33

Verify that the neighbor state with the other two PE routers is Full.

Configuring BGP in the Core

Step-by-Step Procedure

1. On the PE routers, configure BGP. Configure the BGP local autonomous system number.

```
user@PE1# set routing-options autonomous-system 65000
```

```
user@PE2# set routing-options autonomous-system 65000
```

```
user@PE3# set routing-options autonomous-system 65000
```

2. Configure the BGP peer groups. Configure the local address as the `100.0` address on the router. The neighbor addresses are the `100.0` addresses of the other PE routers.

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@PE1# set protocols bgp group group-mvpn type internal
user@PE1# set protocols bgp group group-mvpn local-address 192.168.1.1
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 192.168.2.1
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.7.1
user@PE2# set protocols bgp group group-mvpn type internal
user@PE2# set protocols bgp group group-mvpn local-address 192.168.2.1
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 192.168.1.1
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.7.1
user@PE3# set protocols bgp group group-mvpn type internal
user@PE3# set protocols bgp group group-mvpn local-address 192.168.7.1
user@PE3# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE3# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE3# set protocols bgp group group-mvpn neighbor 192.168.1.1
user@PE3# set protocols bgp group group-mvpn neighbor 192.168.2.1
```

3. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

4. On the PE routers, verify that the BGP neighbors form a peer session.

```

user@PE1> show bgp group
Group Type: Internal      AS: 65000                Local AS: 65000
Name: group-mvpn         Index: 0                  Flags: Export Eval
Holdtime: 0
Total peers: 2           Established: 2
192.168.2.1+54883
192.168.7.1+58933
bgp.l3vpn.0: 0/0/0/0
bgp.mvpn.0: 0/0/0/0

Groups: 1 Peers: 2 External: 0 Internal: 2 Down peers: 0 Flaps: 0
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0      0         0         0         0         0         0
bgp.mvpn.0       0         0         0         0         0         0

```

Verify that the peer state for the other two PE routers is Established and that the `lo0.0` addresses of the other PE routers are shown as peers.

Configuring LDP

Step-by-Step Procedure

1. On the PE routers, configure LDP to support unicast traffic. Specify the core-facing Fast Ethernet and Gigabit Ethernet interfaces between the PE routers. Also configure LDP specifying the `lo0.0` interface. As a best practice, disable LDP on the `fxp0` interface.

```

user@PE1# set protocols ldp deaggregate
user@PE1# set protocols ldp interface fe-0/1/1.0
user@PE1# set protocols ldp interface ge-0/3/0.0
user@PE1# set protocols ldp interface fxp0.0 disable
user@PE1# set protocols ldp interface lo0.0
user@PE2# set protocols ldp deaggregate
user@PE2# set protocols ldp interface fe-0/1/3.0
user@PE2# set protocols ldp interface ge-1/3/0.0
user@PE2# set protocols ldp interface fxp0.0 disable
user@PE2# set protocols ldp interface lo0.0
user@PE3# set protocols ldp deaggregate
user@PE3# set protocols ldp interface fe-0/1/1.0
user@PE3# set protocols ldp interface fe-0/1/3.0

```

```
user@PE3# set protocols ldp interface fxp0.0 disable
user@PE3# set protocols ldp interface lo0.0
```

2. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

3. On the PE routers, use the `show ldp route` command to verify the LDP route.

```
user@PE1> show ldp route
Destination      Next-hop intf/lsp      Next-hop address
10.0.12.8/30     ge-0/3/0.0
10.0.12.9/32
10.0.17.12/30    fe-0/1/1.0
10.0.17.13/32
10.0.27.12/30    fe-0/1/1.0             10.0.17.14
                  ge-0/3/0.0             10.0.12.10
192.168.1.1/32   lo0.0
192.168.2.1/32   ge-0/3/0.0             10.0.12.10
192.168.7.1/32   fe-0/1/1.0             10.0.17.14
224.0.0.5/32
224.0.0.22/32
```

Verify that a next-hop interface and next-hop address have been established for each remote destination in the core network. Notice that local destinations do not have next-hop interfaces, and remote destinations outside the core do not have next-hop addresses.

Configuring RSVP

Step-by-Step Procedure

1. On the PE routers, configure RSVP. Specify the core-facing Fast Ethernet and Gigabit Ethernet interfaces that participate in the LSP. Also specify the `lo0.0` interface. As a best practice, disable RSVP on the `fxp0` interface.

```
user@PE1# set protocols rsvp interface ge-0/3/0.0
user@PE1# set protocols rsvp interface fe-0/1/1.0
user@PE1# set protocols rsvp interface lo0.0
user@PE1# set protocols rsvp interface fxp0.0 disable
user@PE2# set protocols rsvp interface fe-0/1/3.0
user@PE2# set protocols rsvp interface ge-1/3/0.0
user@PE2# set protocols rsvp interface lo0.0
user@PE2# set protocols rsvp interface fxp0.0 disable
user@PE3# set protocols rsvp interface fe-0/1/3.0
user@PE3# set protocols rsvp interface fe-0/1/1.0
user@PE3# set protocols rsvp interface lo0.0
user@PE3# set protocols rsvp interface fxp0.0 disable
```

2. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

Verify these steps using the `show configuration protocols rsvp` command. You can verify the operation of RSVP only after the LSP is established.

Configuring MPLS

Step-by-Step Procedure

1. On the PE routers, configure MPLS. Specify the core-facing Fast Ethernet and Gigabit Ethernet interfaces that participate in the LSP. As a best practice, disable MPLS on the fxp0 interface.

```
user@PE1# set protocols mpls interface ge-0/3/0.0
user@PE1# set protocols mpls interface fe-0/1/1.0
user@PE1# set protocols mpls interface fxp0.0 disable
user@PE2# set protocols mpls interface fe-0/1/3.0
user@PE2# set protocols mpls interface ge-1/3/0.0
user@PE2# set protocols mpls interface fxp0.0 disable
user@PE3# set protocols mpls interface fe-0/1/3.0
user@PE3# set protocols mpls interface fe-0/1/1.0
user@PE3# set protocols mpls interface fxp0.0 disable
```

Use the `show configuration protocols mpls` command to verify that the core-facing Fast Ethernet and Gigabit Ethernet interfaces are configured for MPLS.

2. On the PE routers, configure the core-facing interfaces associated with the LSP. Specify the `mpls` address family type.

```
user@PE1# set interfaces fe-0/1/1 unit 0 family mpls
user@PE1# set interfaces ge-0/3/0 unit 0 family mpls
user@PE2# set interfaces fe-0/1/3 unit 0 family mpls
user@PE2# set interfaces ge-1/3/0 unit 0 family mpls
user@PE3# set interfaces fe-0/1/3 unit 0 family mpls
user@PE3# set interfaces fe-0/1/1 unit 0 family mpls
```

Use the `show mpls interface` command to verify that the core-facing interfaces have the MPLS address family configured.

3. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

You can verify the operation of MPLS after the LSP is established.

Configuring the VRF Routing Instances

Step-by-Step Procedure

1. On Router PE1 , configure the routing instance for the green and red VPNs. Specify the vrf instance type and specify the customer-facing SONET interfaces.

Configure a virtual tunnel (VT) interface on all MVPN routing instances on each PE where hosts in different instances need to receive multicast traffic from the same source.

```
user@PE1# set routing-instances green instance-type vrf
user@PE1# set routing-instances green interface so-0/0/3.0
user@PE1# set routing-instances green interface vt-1/2/0.1 multicast
user@PE1# set routing-instances green interface lo0.1
user@PE1# set routing-instances red instance-type vrf
user@PE1# set routing-instances red interface fe-0/1/0.0
user@PE1# set routing-instances red interface vt-1/2/0.2
user@PE1# set routing-instances red interface lo0.2
```

Use the `show configuration routing-instances green` and `show configuration routing-instances red` commands to verify that the virtual tunnel interfaces have been correctly configured.

2. On Router PE2 , configure the routing instance for the green VPN. Specify the vrf instance type and specify the customer-facing SONET interfaces.

```
user@PE2# set routing-instances green instance-type vrf
user@PE2# set routing-instances green interface so-0/0/1.0
user@PE2# set routing-instances green interface vt-1/2/0.1
user@PE2# set routing-instances green interface lo0.1
```

Use the `show configuration routing-instances green` command.

3. On Router PE3, configure the routing instance for the blue VPN. Specify the vrf instance type and specify the customer-facing SONET interfaces.

```
user@PE3# set routing-instances blue instance-type vrf
user@PE3# set routing-instances blue interface so-0/0/1.0
user@PE3# set routing-instances blue interface vt-1/2/0.3
user@PE3# set routing-instances blue interface lo0.1
```

Use the `show configuration routing-instances blue` command to verify that the instance type has been configured correctly and that the correct interfaces have been configured in the routing instance.

4. On Router PE1, configure a route distinguisher for the green and red routing instances. A route distinguisher allows the router to distinguish between two identical IP prefixes used as VPN routes.



TIP: To help in troubleshooting, this example shows how to configure the route distinguisher to match the router ID. This allows you to associate a route with the router that advertised it.

```
user@PE1# set routing-instances green route-distinguisher 192.168.1.1:1
user@PE1# set routing-instances red route-distinguisher 192.168.1.1:2
```

5. On Router PE2, configure a route distinguisher for the green routing instance.

```
user@PE2# set routing-instances green route-distinguisher 192.168.2.1:1
```

6. On Router PE3, configure a route distinguisher for the blue routing instance.

```
user@PE3# set routing-instances blue route-distinguisher 192.168.7.1:3
```

7. On the PE routers, configure the VPN routing instance for multicast support.

```
user@PE1# set routing-instances green protocols mvpn
user@PE1# set routing-instances red protocols mvpn
user@PE2# set routing-instances green protocols mvpn
user@PE3# set routing-instances blue protocols mvpn
```

Use the `show configuration routing-instance` command to verify that the route distinguisher is configured correctly and that the MVPN Protocol is enabled in the routing instance.

8. On the PE routers, configure an IP address on additional loopback logical interfaces. These logical interfaces are used as the loopback addresses for the VPNs.

```
user@PE1# set interfaces lo0 unit 1 description "green VRF loopback"
user@PE1# set interfaces lo0 unit 1 family inet address 10.10.1.1/32
user@PE1# set interfaces lo0 unit 2 description "red VRF loopback"
user@PE1# set interfaces lo0 unit 2 family inet address 10.2.1.1/32
user@PE2# set interfaces lo0 unit 1 description "green VRF loopback"
user@PE2# set interfaces lo0 unit 1 family inet address 10.10.22.2/32
user@PE3# set interfaces lo0 unit 1 description "blue VRF loopback"
user@PE3# set interfaces lo0 unit 1 family inet address 10.3.33.3/32
```

Use the `show interfaces terse` command to verify that the loopback logical interfaces are correctly configured.

9. On the PE routers, configure virtual tunnel interfaces. These interfaces are used in VRF instances where multicast traffic arriving on a provider tunnel needs to be forwarded to multiple VPNs.

```
user@PE1# set interfaces vt-1/2/0 unit 1 description "green VRF multicast vt"
user@PE1# set interfaces vt-1/2/0 unit 1 family inet
user@PE1# set interfaces vt-1/2/0 unit 2 description "red VRF unicast and multicast vt"
user@PE1# set interfaces vt-1/2/0 unit 2 family inet
user@PE1# set interfaces vt-1/2/0 unit 3 description "blue VRF multicast vt"
user@PE1# set interfaces vt-1/2/0 unit 3 family inet
user@PE2# set interfaces vt-1/2/0 unit 1 description "green VRF unicast and multicast vt"
user@PE2# set interfaces vt-1/2/0 unit 1 family inet
```

```

user@PE2# set interfaces vt-1/2/0 unit 3 description "blue VRF unicast and multicast vt"
user@PE2# set interfaces vt-1/2/0 unit 3 family inet
user@PE3# set interfaces vt-1/2/0 unit 3 description "blue VRF unicast and multicast vt"
user@PE3# set interfaces vt-1/2/0 unit 3 family inet

```

Use the `show interfaces terse` command to verify that the virtual tunnel interfaces have the correct address family type configured.

10. On the PE routers, configure the provider tunnel.

```

user@PE1# set routing-instances green provider-tunnel rsvp-te label-switched-path-template
default-template
user@PE1# set routing-instances red provider-tunnel rsvp-te label-switched-path-template
default-template
user@PE2# set routing-instances green provider-tunnel rsvp-te label-switched-path-template
default-template
user@PE3# set routing-instances blue provider-tunnel rsvp-te label-switched-path-template
default-template

```

Use the `show configuration routing-instance` command to verify that the provider tunnel is configured to use the default LSP template.



NOTE: You cannot commit the configuration for the VRF instance until you configure the VRF target in the next section.

Configuring MVPN Extranet Policy

Step-by-Step Procedure

1. On the PE routers, define the VPN community name for the route targets for each VPN. The community names are used in the VPN import and export policies.

```

user@PE1# set policy-options community green-com members target:65000:1
user@PE1# set policy-options community red-com members target:65000:2
user@PE1# set policy-options community blue-com members target:65000:3
user@PE2# set policy-options community green-com members target:65000:1
user@PE2# set policy-options community red-com members target:65000:2
user@PE2# set policy-options community blue-com members target:65000:3
user@PE3# set policy-options community green-com members target:65000:1

```

```

user@PE3# set policy-options community red-com members target:65000:2
user@PE3# set policy-options community blue-com members target:65000:3

```

Use the `show policy-options` command to verify that the correct VPN community name and route target are configured.

2. On the PE routers, configure the VPN import policy. Include the community name of the route targets that you want to accept. Do not include the community name of the route targets that you do not want to accept. For example, omit the community name for routes from the VPN of a multicast sender from which you do not want to receive multicast traffic.

```

user@PE1# set policy-options policy-statement green-red-blue-import term t1 from community
green-com
user@PE1# set policy-options policy-statement green-red-blue-import term t1 from community
red-com
user@PE1# set policy-options policy-statement green-red-blue-import term t1 from community
blue-com
user@PE1# set policy-options policy-statement green-red-blue-import term t1 then accept
user@PE1# set policy-options policy-statement green-red-blue-import term t2 then reject
user@PE2# set policy-options policy-statement green-red-blue-import term t1 from community
green-com
user@PE2# set policy-options policy-statement green-red-blue-import term t1 from community
red-com
user@PE2# set policy-options policy-statement green-red-blue-import term t1 from community
blue-com
user@PE2# set policy-options policy-statement green-red-blue-import term t1 then accept
user@PE2# set policy-options policy-statement green-red-blue-import term t2 then reject
user@PE3# set policy-options policy-statement green-red-blue-import term t1 from community
green-com
user@PE3# set policy-options policy-statement green-red-blue-import term t1 from community
red-com
user@PE3# set policy-options policy-statement green-red-blue-import term t1 from community
blue-com
user@PE3# set policy-options policy-statement green-red-blue-import term t1 then accept
user@PE3# set policy-options policy-statement green-red-blue-import term t2 then reject

```

Use the `show policy green-red-blue-import` command to verify that the VPN import policy is correctly configured.

3. On the PE routers, apply the VRF import policy. In this example, the policy is defined in a policy-statement policy, and target communities are defined under the [edit policy-options] hierarchy level.

```
user@PE1# set routing-instances green vrf-import green-red-blue-import
user@PE1# set routing-instances red vrf-import green-red-blue-import
user@PE2# set routing-instances green vrf-import green-red-blue-import
user@PE3# set routing-instances blue vrf-import green-red-blue-import
```

Use the `show configuration routing-instances` command to verify that the correct VRF import policy has been applied.

4. On the PE routers, configure VRF export targets. The `vrf-target` statement and `export` option cause the routes being advertised to be labeled with the target community.

For Router PE3, the `vrf-target` statement is included without specifying the `export` option. If you do not specify the `import` or `export` options, default VRF import and export policies are generated that accept imported routes and tag exported routes with the specified target community.



NOTE: You must configure the same route target on each PE router for a given VPN routing instance.

```
user@PE1# set routing-instances green vrf-target export target:65000:1
user@PE1# set routing-instances red vrf-target export target:65000:2
user@PE2# set routing-instances green vrf-target export target:65000:1
user@PE3# set routing-instances blue vrf-target target:65000:3
```

Use the `show configuration routing-instances` command to verify that the correct VRF export targets have been configured.

5. On the PE routers, configure automatic exporting of routes between VRF instances. When you include the `auto-export` statement, the `vrf-import` and `vrf-export` policies are compared across all VRF instances. If there is a common route target community between the instances, the routes are shared. In this example, the `auto-export` statement must be included under all instances that need to send traffic to and receive traffic from another instance located on the same router.

```
user@PE1# set routing-instances green routing-options auto-export
user@PE1# set routing-instances red routing-options auto-export
user@PE2# set routing-instances green routing-options auto-export
user@PE3# set routing-instances blue routing-options auto-export
```


- On the PE routers, configure the load balance policy statement. While load balancing leads to better utilization of the available links, it is not required for MVPN extranets. It is included here as a best practice.

```

user@PE1# set policy-options policy-statement load-balance then load-balance per-packet
user@PE2# set policy-options policy-statement load-balance then load-balance per-packet
user@PE3# set policy-options policy-statement load-balance then load-balance per-packet

```

Use the `show policy-options` command to verify that the load balance policy statement has been correctly configured.

- On the PE routers, apply the load balance policy.

```

user@PE1# set routing-options forwarding-table export load-balance
user@PE2# set routing-options forwarding-table export load-balance
user@PE3# set routing-options forwarding-table export load-balance

```

- On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

- On the PE routers, use the `show rsvp neighbor` command to verify that the RSVP neighbors are established.

```

user@PE1> show rsvp neighbor
RSVP neighbor: 2 learned
Address           Idle Up/Dn  LastChange HelloInt HelloTx/Rx MsgRcvd
10.0.17.14        5  1/0    43:52     9   293/293  247
10.0.12.10        0  1/0    50:15     9   336/336  140

```

Verify that the other PE routers are listed as RSVP neighbors.

10. On the PE routers, display the MPLS LSPs.

```

user@PE1> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: 192.168.1.1:1:mvpn:green, P2MP branch count: 2
To          From          State Rt P    ActivePath      LSPname
192.168.2.1 192.168.1.1    Up    0 *
192.168.2.1:192.168.1.1:1:mvpn:green
192.168.7.1 192.168.1.1    Up    0 *
192.168.7.1:192.168.1.1:1:mvpn:green
P2MP name: 192.168.1.1:2:mvpn:red, P2MP branch count: 2
To          From          State Rt P    ActivePath      LSPname
192.168.2.1 192.168.1.1    Up    0 *
192.168.2.1:192.168.1.1:2:mvpn:red
192.168.7.1 192.168.1.1    Up    0 *
192.168.7.1:192.168.1.1:2:mvpn:red
Total 4 displayed, Up 4, Down 0

Egress LSP: 2 sessions
P2MP name: 192.168.2.1:1:mvpn:green, P2MP branch count: 1
To          From          State  Rt Style Labelin Labelout LSPname
192.168.1.1 192.168.2.1    Up     0  1 SE 299888      3
192.168.1.1:192.168.2.1:1:mvpn:green
P2MP name: 192.168.7.1:3:mvpn:blue, P2MP branch count: 1
To          From          State  Rt Style Labelin Labelout LSPname
192.168.1.1 192.168.7.1    Up     0  1 SE 299872      3
192.168.1.1:192.168.7.1:3:mvpn:blue
Total 2 displayed, Up 2, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

In this display from Router PE1, notice that there are two ingress LSPs for the green VPN and two for the red VPN configured on this router. Verify that the state of each ingress LSP is up. Also notice that there is one egress LSP for each of the green and blue VPNs. Verify that the state of each egress LSP is up.



TIP: The LSP name displayed in the `show mpls lsp p2mp` command output can be used in the `ping mpls rsvp <lsp-name> multipath` command.

Configuring CE-PE BGP

Step-by-Step Procedure

1. On the PE routers, configure the BGP export policy. The BGP export policy is used to allow static routes and routes that originated from directly attached interfaces to be exported to BGP.

```

user@PE1# set policy-options policy-statement BGP-export term t1 from protocol direct
user@PE1# set policy-options policy-statement BGP-export term t1 then accept
user@PE1# set policy-options policy-statement BGP-export term t2 from protocol static
user@PE1# set policy-options policy-statement BGP-export term t2 then accept
user@PE2# set policy-options policy-statement BGP-export term t1 from protocol direct
user@PE2# set policy-options policy-statement BGP-export term t1 then accept
user@PE2# set policy-options policy-statement BGP-export term t2 from protocol static
user@PE2# set policy-options policy-statement BGP-export term t2 then accept
user@PE3# set policy-options policy-statement BGP-export term t1 from protocol direct
user@PE3# set policy-options policy-statement BGP-export term t1 then accept
user@PE3# set policy-options policy-statement BGP-export term t2 from protocol static
user@PE3# set policy-options policy-statement BGP-export term t2 then accept

```

Use the `show policy BGP-export` command to verify that the BGP export policy is correctly configured.

2. On the PE routers, configure the CE to PE BGP session. Use the IP address of the SONET interface as the neighbor address. Specify the autonomous system number for the VPN network of the attached CE router.

```

user@PE1# set routing-instances green protocols bgp group PE-CE export BGP-export
user@PE1# set routing-instances green protocols bgp group PE-CE neighbor 10.0.16.1 peer-as
65001
user@PE2# set routing-instances green protocols bgp group PE-CE export BGP-export
user@PE2# set routing-instances green protocols bgp group PE-CE neighbor 10.0.24.2 peer-as
65009
user@PE3# set routing-instances blue protocols bgp group PE-CE export BGP-export
user@PE3# set routing-instances blue protocols bgp group PE-CE neighbor 10.0.79.2 peer-as
65003

```

3. On the CE routers, configure the BGP local autonomous system number.

```
user@CE1# set routing-options autonomous-system 65001
user@CE2# set routing-options autonomous-system 65009
user@CE3# set routing-options autonomous-system 65003
```

4. On the CE routers, configure the BGP export policy. The BGP export policy is used to allow static routes and routes that originated from directly attached interfaces to be exported to BGP.

```
user@CE1# set policy-options policy-statement BGP-export term t1 from protocol direct
user@CE1# set policy-options policy-statement BGP-export term t1 then accept
user@CE1# set policy-options policy-statement BGP-export term t2 from protocol static
user@CE1# set policy-options policy-statement BGP-export term t2 then accept
user@CE2# set policy-options policy-statement BGP-export term t1 from protocol direct
user@CE2# set policy-options policy-statement BGP-export term t1 then accept
user@CE2# set policy-options policy-statement BGP-export term t2 from protocol static
user@CE2# set policy-options policy-statement BGP-export term t2 then accept
user@CE3# set policy-options policy-statement BGP-export term t1 from protocol direct
user@CE3# set policy-options policy-statement BGP-export term t1 then accept
user@CE3# set policy-options policy-statement BGP-export term t2 from protocol static
user@CE3# set policy-options policy-statement BGP-export term t2 then accept
```

Use the `show policy BGP-export` command to verify that the BGP export policy is correctly configured.

5. On the CE routers, configure the CE-to-PE BGP session. Use the IP address of the SONET interface as the neighbor address. Specify the autonomous system number of the core network. Apply the BGP export policy.

```
user@CE1# set protocols bgp group PE-CE export BGP-export
user@CE1# set protocols bgp group PE-CE neighbor 10.0.16.2 peer-as 65000
user@CE2# set protocols bgp group PE-CE export BGP-export
user@CE2# set protocols bgp group PE-CE neighbor 10.0.24.1 peer-as 65000
user@CE3# set protocols bgp group PE-CE export BGP-export
user@CE3# set protocols bgp group PE-CE neighbor 10.0.79.1 peer-as 65000
```

6. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

7. On the PE routers, use the `show bgp group pe-ce` command to verify that the BGP neighbors form a peer session.

```
user@PE1> show bgp group pe-ce
Group Type: External                               Local AS: 65000
Name: PE-CE           Index: 1                     Flags: <>
Export: [ BGP-export ]
Holdtime: 0
Total peers: 1      Established: 1
10.0.16.1+60500
green.inet.0: 2/3/3/0
```

Verify that the peer state for the CE routers is `Established` and that the IP address configured on the peer SONET interface is shown as the peer.

Configuring PIM on the PE Routers

Step-by-Step Procedure

1. On the PE routers, enable an instance of PIM in each VPN. Configure the `lo0.1`, `lo0.2`, and customer-facing SONET and Fast Ethernet interfaces. Specify the mode as `sparse`.

```
user@PE1# set routing-instances green protocols pim interface lo0.1 mode sparse
user@PE1# set routing-instances green protocols pim interface so-0/0/3.0 mode sparse
user@PE1# set routing-instances red protocols pim interface lo0.2 mode sparse
user@PE1# set routing-instances red protocols pim interface fe-0/1/0.0 mode sparse
```

```

user@PE2# set routing-instances green protocols pim interface lo0.1 mode sparse
user@PE2# set routing-instances green protocols pim interface so-0/0/1.0 mode sparse
user@PE3# set routing-instances blue protocols pim interface lo0.1 mode sparse
user@PE3# set routing-instances blue protocols pim interface so-0/0/1.0 mode sparse

```

2. On the PE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

3. On the PE routers, use the `show pim interfaces instance green` command and substitute the appropriate VRF instance name to verify that the PIM interfaces are up.

```

user@PE1> show pim interfaces instance green
Instance: PIM.green

Name           Stat Mode      IP V State NbrCnt JoinCnt DR address
lo0.1          Up   Sparse       4 2 DR      0      0 10.10.1.1
lsi.0          Up   SparseDense  4 2 P2P     0      0
pe-1/2/0.32769 Up   Sparse       4 2 P2P     0      0
so-0/0/3.0     Up   Sparse       4 2 P2P     1      2
vt-1/2/0.1    Up   SparseDense  4 2 P2P     0      0
lsi.0          Up   SparseDense  6 2 P2P     0      0

```

Also notice that the normal mode for the virtual tunnel interface and label-switched interface is SparseDense.

Configuring PIM on the CE Routers

Step-by-Step Procedure

1. On the CE routers, configure the customer-facing and core-facing interfaces for PIM. Specify the mode as sparse.

```

user@CE1# set protocols pim interface fe-1/3/0.0 mode sparse
user@CE1# set protocols pim interface so-0/0/3.0 mode sparse
user@CE2# set protocols pim interface fe-0/1/1.0 mode sparse
user@CE2# set protocols pim interface so-0/0/1.0 mode sparse
user@CE3# set protocols pim interface fe-0/1/0.0 mode sparse
user@CE3# set protocols pim interface so-0/0/1.0 mode sparse

```

Use the `show pim interfaces` command to verify that the PIM interfaces have been configured to use sparse mode.

2. On the CE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

3. On the CE routers, use the `show pim interfaces` command to verify that the PIM interface status is up.

```
user@CE1> show pim interfaces
```

```
Instance: PIM.master
```

Name	Stat	Mode	IP V State	NbrCnt	JoinCnt	DR address
fe-1/3/0.0	Up	Sparse	4 2 DR	0	0	10.10.12.1
pe-1/2/0.32769	Up	Sparse	4 2 P2P	0	0	
so-0/0/3.0	Up	Sparse	4 2 P2P	1	1	

Configuring the Rendezvous Points

Step-by-Step Procedure

1. Configure Router PE1 to be the rendezvous point for the red VPN instance of PIM. Specify the local 100.2 address.

```
user@PE1# set routing-instances red protocols pim rp local address 10.2.1.1
```

2. Configure Router PE2 to be the rendezvous point for the green VPN instance of PIM. Specify the 100.1 address of Router PE2.

```
user@PE2# set routing-instances green protocols pim rp local address 10.10.22.2
```

3. Configure Router PE3 to be the rendezvous point for the blue VPN instance of PIM. Specify the local 100.1.

```
user@PE3# set routing-instances blue protocols pim rp local address 10.3.33.3
```

4. On the PE1, CE1, and CE2 routers, configure the static rendezvous point for the green VPN instance of PIM. Specify the 100.1 address of Router PE2.

```
user@PE1# set routing-instances green protocols pim rp static address 10.10.22.2  
user@CE1# set protocols pim rp static address 10.10.22.2  
user@CE2# set protocols pim rp static address 10.10.22.2
```

5. On Router CE3, configure the static rendezvous point for the blue VPN instance of PIM. Specify the 100.1 address of Router PE3.

```
user@CE3# set protocols pim rp static address 10.3.33.3
```


6. On the CE routers, commit the configuration:

```
user@host> commit check
```

```
configuration check succeeds
```

```
user@host> commit
```

```
commit complete
```

7. On the PE routers, use the `show pim rps instance <instance-name>` command and substitute the appropriate VRF instance name to verify that the RPs have been correctly configured.

```
user@PE1> show pim rps instance <instance-name>
Instance: PIM.green
Address family INET
RP address          Type          Holdtime Timeout Groups Group prefixes
10.10.22.2          static        0      None    1 224.0.0.0/4

Address family INET6
```

Verify that the correct IP address is shown as the RP.

8. On the CE routers, use the `show pim rps` command to verify that the RP has been correctly configured.

```
user@CE1> show pim rps
Instance: PIM.master
Address family INET
RP address          Type          Holdtime Timeout Groups Group prefixes
10.10.22.2          static        0      None    1 224.0.0.0/4

Address family INET6
```

Verify that the correct IP address is shown as the RP.

9. On Router PE1, use the `show route table green.mvpn.0 | find 1` command to verify that the type-1 routes have been received from the PE2 and PE3 routers.

```

user@PE1> show route table green.mvpn.0 | find
1
green.mvpn.0: 7 destinations, 9 routes (7 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:192.168.1.1:1:192.168.1.1/240
          *[MVPN/70] 03:38:09, metric2 1
          Indirect
1:192.168.1.1:2:192.168.1.1/240
          *[MVPN/70] 03:38:05, metric2 1
          Indirect
1:192.168.2.1:1:192.168.2.1/240
          *[BGP/170] 03:12:18, localpref 100, from 192.168.2.1
          AS path: I
          > to 10.0.12.10 via ge-0/3/0.0
1:192.168.7.1:3:192.168.7.1/240
          *[BGP/170] 03:12:18, localpref 100, from 192.168.7.1
          AS path: I
          > to 10.0.17.14 via fe-0/1/1.0

```

10. On Router PE1, use the `show route table green.mvpn.0 | find 5` command to verify that the type-5 routes have been received from Router PE2.

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a PIM router learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface. If an MBGP MVPN is also configured, the PE device originates a type-5 MVPN route.

```

user@PE1> show route table green.mvpn.0 | find
5
5:192.168.2.1:1:32:10.10.12.52:32:224.1.1.1/240
          *[BGP/170] 03:12:18, localpref 100, from 192.168.2.1
          AS path: I
          > to 10.0.12.10 via ge-0/3/0.0

```

11. On Router PE1, use the `show route table green.mvpn.0 | find 7` command to verify that the type-7 routes have been received from Router PE2.

```

user@PE1> show route table green.mvpn.0 | find
7
7:192.168.1.1:1:65000:32:10.10.12.52:32:224.1.1.1/240
    *[MVPN/70] 03:22:47, metric2 1
        Multicast (IPv4)
    [PIM/105] 03:34:18
        Multicast (IPv4)
    [BGP/170] 03:12:18, localpref 100, from 192.168.2.1
        AS path: I
        > to 10.0.12.10 via ge-0/3/0.0

```

12. On Router PE1, use the `show route advertising-protocol bgp 192.168.2.1 table green.mvpn.0 detail` command to verify that the routes advertised by Router PE2 use the PMSI attribute set to RSVP-TE.

```

user@PE1> show route advertising-protocol bgp
192.168.2.1 table green.mvpn.0 detail
green.mvpn.0: 7 destinations, 9 routes (7 active, 1 holddown, 0 hidden)
* 1:192.168.1.1:1:192.168.1.1/240 (1 entry, 1 announced)
  BGP group group-mvpn type Internal
    Route Distinguisher: 192.168.1.1:1
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:65000:1
    PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[192.168.1.1:0:56822:192.168.1.1]

```

Testing MVPN Extranets

Step-by-Step Procedure

1. Start the multicast receiver device connected to Router CE2.
2. Start the multicast sender device connected to Router CE1.
3. Verify that the receiver receives the multicast stream.

4. On Router PE1, display the provider tunnel to multicast group mapping by using the `show mvpn c-multicast` command.

```

user@PE1> 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: green
  C-mcast IPv4 (S:G)          Ptnl          St
  10.10.12.52/32:224.1.1.1/32  RSVP-TE P2MP:192.168.1.1, 56822,192.168.1.1  RM
  0.0.0.0/0:239.255.255.250/32
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: red
  C-mcast IPv4 (S:G)          Ptnl          St          DS
  10.10.12.52/32:224.1.1.1/32                                     DS
  0.0.0.0/0:224.1.1.1/32

```

5. On Router PE2, use the `show route table green.mvpn.0 | find 6` command to verify that the type-6 routes have been created as a result of receiving PIM join messages.

```

user@PE2> show route table green.mvpn.0 | find
6
6:192.168.2.1:1:65000:32:10.10.22.2:32:224.1.1.1/240
    *[PIM/105] 04:01:23
    Multicast (IPv4)
6:192.168.2.1:1:65000:32:10.10.22.2:32:239.255.255.250/240
    *[PIM/105] 22:39:46
    Multicast (IPv4)

```



NOTE: The multicast address 239.255.255.250 shown in the preceding step is not related to this example. This address is sent by some host machines.

6. Start the multicast receiver device connected to Router CE3.
7. Verify that the receiver is receiving the multicast stream.
8. On Router PE2, use the `show route table green.mvpn.0 | find 6` command to verify that the type-6 routes have been created as a result of receiving PIM join messages from the multicast receiver device connected to Router CE3.

```
user@PE2> show route table green.mvpn.0 | find
6
6:192.168.2.1:1:65000:32:10.10.22.2:32:239.255.255.250/240
      *[PIM/105] 06:43:39
      Multicast (IPv4)
```

9. Start the multicast receiver device directly connected to Router PE1.
10. Verify that the receiver is receiving the multicast stream.
11. On Router PE1, use the `show route table green.mvpn.0 | find 6` command to verify that the type-6 routes have been created as a result of receiving PIM join messages from the directly connected multicast receiver device.

```
user@PE1> show route table green.mvpn.0 | find
6
6:192.168.1.1:2:65000:32:10.2.1.1:32:224.1.1.1/240
      *[PIM/105] 00:02:32
      Multicast (IPv4)
6:192.168.1.1:2:65000:32:10.2.1.1:32:239.255.255.250/240
      *[PIM/105] 00:05:49
      Multicast (IPv4)
```



NOTE: The multicast address 255.255.255.250 shown in the step above is not related to this example.

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 {
  so-0/0/3 {
    unit 0 {
      description "to PE1 so-0/0/3.0";
      family inet {
        address 10.0.16.1/30;
      }
    }
  }
  fe-1/3/0 {
    unit 0 {
      family inet {
        address 10.10.12.1/24;
      }
    }
  }
  lo0 {
    unit 0 {
      description "CE1 Loopback";
      family inet {
        address 192.168.6.1/32 {
          primary;
        }
        address 127.0.0.1/32;
      }
    }
  }
}
routing-options {
  autonomous-system 65001;
  router-id 192.168.6.1;
  forwarding-table {
    export load-balance;
  }
}
```

```
}
protocols {
  bgp {
    group PE-CE {
      export BGP-export;
      neighbor 10.0.16.2 {
        peer-as 65000;
      }
    }
  }
  pim {
    rp {
      static {
        address 10.10.22.2;
      }
    }
    interface fe-1/3/0.0 {
      mode sparse;
    }
    interface so-0/0/3.0 {
      mode sparse;
    }
  }
}
policy-options {
  policy-statement BGP-export {
    term t1 {
      from protocol direct;
      then accept;
    }
    term t2 {
      from protocol static;
      then accept;
    }
  }
  policy-statement load-balance {
    then {
      load-balance per-packet;
    }
  }
}
```

The relevant sample configuration for Router PE1 follows.

Router PE1

```
interfaces {
  so-0/0/3 {
    unit 0 {
      description "to CE1 so-0/0/3.0";
      family inet {
        address 10.0.16.2/30;
      }
    }
  }
  fe-0/1/0 {
    unit 0 {
      description "to H2";
      family inet {
        address 10.2.11.2/30;
      }
    }
  }
  fe-0/1/1 {
    unit 0 {
      description "to PE3 fe-0/1/1.0";
      family inet {
        address 10.0.17.13/30;
      }
      family mpls;
    }
  }
  ge-0/3/0 {
    unit 0 {
      description "to PE2 ge-1/3/0.0";
      family inet {
        address 10.0.12.9/30;
      }
      family mpls;
    }
  }
  vt-1/2/0 {
    unit 1 {
      description "green VRF multicast vt";
      family inet;
    }
  }
}
```



```
    unit 2 {
      description "red VRF unicast and multicast vt";
      family inet;
    }
    unit 3 {
      description "blue VRF multicast vt";
      family inet;
    }
  }
  lo0 {
    unit 0 {
      description "PE1 Loopback";
      family inet {
        address 192.168.1.1/32 {
          primary;
        }
        address 127.0.0.1/32;
      }
    }
    unit 1 {
      description "green VRF loopback";
      family inet {
        address 10.10.1.1/32;
      }
    }
    unit 2 {
      description "red VRF loopback";
      family inet {
        address 10.2.1.1/32;
      }
    }
  }
}
routing-options {
  autonomous-system 65000;
  router-id 192.168.1.1;
  forwarding-table {
    export load-balance;
  }
}
protocols {
  rsvp {
    interface ge-0/3/0.0;
```

```
interface fe-0/1/1.0;
interface lo0.0;
interface fxp0.0 {
    disable;
}
}
mpls {
    interface ge-0/3/0.0;
    interface fe-0/1/1.0;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group group-mvpn {
        type internal;
        local-address 192.168.1.1;
        family inet-vpn {
            unicast;
        }
        family inet-mvpn {
            signaling;
        }
        neighbor 192.168.2.1;
        neighbor 192.168.7.1;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/3/0.0 {
            metric 100;
        }
        interface fe-0/1/1.0 {
            metric 100;
        }
        interface lo0.0 {
            passive;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

```
}
ldp {
    deaggregate;
    interface ge-0/3/0.0;
    interface fe-0/1/1.0;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
}
policy-options {
    policy-statement BGP-export {
        term t1 {
            from protocol direct;
            then accept;
        }
        term t2 {
            from protocol static;
            then accept;
        }
    }
    policy-statement green-red-blue-import {
        term t1 {
            from community [ green-com red-com blue-com ];
            then accept;
        }
        term t2 {
            then reject;
        }
    }
    policy-statement load-balance {
        then {
            load-balance per-packet;
        }
    }
    community green-com members target:65000:1;
    community red-com members target:65000:2;
    community blue-com members target:65000:3;
}
routing-instances {
    green {
        instance-type vrf;
    }
}
```

```
interface so-0/0/3.0;
interface vt-1/2/0.1 {
    multicast;
}
interface lo0.1;
route-distinguisher 192.168.1.1:1;
provider-tunnel {
    rsvp-te {
        label-switched-path-template {
            default-template;
        }
    }
}
vrf-import green-red-blue-import;
vrf-target export target:65000:1;
vrf-table-label;
routing-options {
    auto-export;
}
protocols {
    bgp {
        group PE-CE {
            export BGP-export;
            neighbor 10.0.16.1 {
                peer-as 65001;
            }
        }
    }
    pim {
        rp {
            static {
                address 10.10.22.2;
            }
        }
        interface so-0/0/3.0 {
            mode sparse;
        }
        interface lo0.1 {a
            mode sparse;
        }
    }
}
mvpn;
}
```

```

red {
  instance-type vrf;
  interface fe-0/1/0.0;
  interface vt-1/2/0.2;
  interface lo0.2;
  route-distinguisher 192.168.1.1:2;
  provider-tunnel {
    rsvp-te {
      label-switched-path-template {
        default-template;
      }
    }
  }
  vrf-import green-red-blue-import;
  vrf-target export target:65000:2;
  routing-options {
    auto-export;
  }
  protocols {
    pim {
      rp {
        local {
          address 10.2.1.1;
        }
      }
      interface fe-0/1/0.0 {
        mode sparse;
      }
      interface lo0.2 {
        mode sparse;
      }
    }
    mvpn;
  }
}
}

```

The relevant sample configuration for Router PE2 follows.

Router PE2

```
interfaces {
  so-0/0/1 {
    unit 0 {
      description "to CE2 so-0/0/1:0.0";
      family inet {
        address 10.0.24.1/30;
      }
    }
  }
  fe-0/1/3 {
    unit 0 {
      description "to PE3 fe-0/1/3.0";
      family inet {
        address 10.0.27.13/30;
      }
      family mpls;
    }
  }
  vt-1/2/0 {
    unit 1 {
      description "green VRF unicast and multicast vt";
      family inet;
    }
    unit 3 {
      description "blue VRF unicast and multicast vt";
      family inet;
    }
  }
}
ge-1/3/0 {
  unit 0 {
    description "to PE1 ge-0/3/0.0";
    family inet {
      address 10.0.12.10/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    description "PE2 Loopback";
```

```
    family inet {
        address 192.168.2.1/32 {
            primary;
        }
        address 127.0.0.1/32;
    }
}
unit 1 {
    description "green VRF loopback";
    family inet {
        address 10.10.22.2/32;
    }
}
}
routing-options {
    router-id 192.168.2.1;
    autonomous-system 65000;
    forwarding-table {
        export load-balance;
    }
}
}
protocols {
    rsvp {
        interface fe-0/1/3.0;
        interface ge-1/3/0.0;
        interface lo0.0;
        interface fxp0.0 {
            disable;
        }
    }
}
mpls {
    interface fe-0/1/3.0;
    interface ge-1/3/0.0;
    interface fxp0.0 {
        disable;
    }
}
}
bgp {
    group group-mvpn {
        type internal;
        local-address 192.168.2.1;
        family inet-vpn {
            unicast;
        }
    }
}
```

```
    }
    family inet-mvpn {
        signaling;
    }
    neighbor 192.168.1.1;
    neighbor 192.168.7.1;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-0/1/3.0 {
            metric 100;
        }
        interface ge-1/3/0.0 {
            metric 100;
        }
        interface lo0.0 {
            passive;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    deaggregate;
    interface fe-0/1/3.0;
    interface ge-1/3/0.0;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
}
policy-options {
    policy-statement BGP-export {
        term t1 {
            from protocol direct;
            then accept;
        }
        term t2 {
            from protocol static;
```



```

        then accept;
    }
}
policy-statement green-red-blue-import {
    term t1 {
        from community [ green-com red-com blue-com ];
        then accept;
    }
    term t2 {
        then reject;
    }
}
policy-statement load-balance {
    then {
        load-balance per-packet;
    }
}
community green-com members target:65000:1;
community red-com members target:65000:2;
community blue-com members target:65000:3;
}
routing-instances {
    green {
        instance-type vrf;
        interface so-0/0/1.0;
        interface vt-1/2/0.1;
        interface lo0.1;
        route-distinguisher 192.168.2.1:1;
        provider-tunnel {
            rsvp-te {
                label-switched-path-template {
                    default-template;
                }
            }
        }
        vrf-import green-red-blue-import;
        vrf-target export target:65000:1;
        routing-options {
            auto-export;
        }
        protocols {
            bgp {
                group PE-CE {

```



```
        address 10.0.24.2/30;
    }
}
lo0 {
    unit 0 {
        description "CE2 Loopback";
        family inet {
            address 192.168.4.1/32 {
                primary;
            }
            address 127.0.0.1/32;
        }
    }
}
routing-options {
    router-id 192.168.4.1;
    autonomous-system 65009;
    forwarding-table {
        export load-balance;
    }
}
protocols {
    bgp {
        group PE-CE {
            export BGP-export;
            neighbor 10.0.24.1 {
                peer-as 65000;
            }
        }
    }
    pim {
        rp {
            static {
                address 10.10.22.2;
            }
        }
        interface so-0/0/1.0 {
            mode sparse;
        }
        interface fe-0/1/1.0 {
            mode sparse;
        }
    }
}
```

```

    }
  }
}
policy-options {
  policy-statement BGP-export {
    term t1 {
      from protocol direct;
      then accept;
    }
    term t2 {
      from protocol static;
      then accept;
    }
  }
  policy-statement load-balance {
    then {
      load-balance per-packet;
    }
  }
}
}

```

The relevant sample configuration for Router PE3 follows.

Router PE3

```

interfaces {
  so-0/0/1 {
    unit 0 {
      description "to CE3 so-0/0/1.0";
      family inet {
        address 10.0.79.1/30;
      }
    }
  }
  fe-0/1/1 {
    unit 0 {
      description "to PE1 fe-0/1/1.0";
      family inet {
        address 10.0.17.14/30;
      }
      family mpls;
    }
  }
}

```

```
}
fe-0/1/3 {
  unit 0 {
    description "to PE2 fe-0/1/3.0";
    family inet {
      address 10.0.27.14/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 3 {
    description "blue VRF unicast and multicast vt";
    family inet;
  }
}
lo0 {
  unit 0 {
    description "PE3 Loopback";
    family inet {
      address 192.168.7.1/32 {
        primary;
      }
      address 127.0.0.1/32;
    }
  }
  unit 1 {
    description "blue VRF loopback";
    family inet {
      address 10.3.33.3/32;
    }
  }
}
}
routing-options {
  router-id 192.168.7.1;
  autonomous-system 65000;
  forwarding-table {
    export load-balance;
  }
}
}
protocols {
  rsvp {
```

```
interface fe-0/1/3.0;
interface fe-0/1/1.0;
interface lo0.0;
interface fxp0.0 {
    disable;
}
}
mpls {
    interface fe-0/1/3.0;
    interface fe-0/1/1.0;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group group-mvpn {
        type internal;
        local-address 192.168.7.1;
        family inet-vpn {
            unicast;
        }
        family inet-mvpn {
            signaling;
        }
        neighbor 192.168.1.1;
        neighbor 192.168.2.1;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-0/1/3.0 {
            metric 100;
        }
        interface fe-0/1/1.0 {
            metric 100;
        }
        interface lo0.0 {
            passive;
        }
        interface fxp0.0 {
            disable;
        }
    }
}
```

```

    }
}
ldp {
    deaggregate;
    interface fe-0/1/3.0;
    interface fe-0/1/1.0;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
}
policy-options {
    policy-statement BGP-export {
        term t1 {
            from protocol direct;
            then accept;
        }
        term t2 {
            from protocol static;
            then accept;
        }
    }
    policy-statement green-red-blue-import {
        term t1 {
            from community [ green-com red-com blue-com ];
            then accept;
        }
        term t2 {
            then reject;
        }
    }
    policy-statement load-balance {
        then {
            load-balance per-packet;
        }
    }
    community green-com members target:65000:1;
    community red-com members target:65000:2;
    community blue-com members target:65000:3;
}
routing-instances {
    blue {

```

```
instance-type vrf;
interface vt-1/2/0.3;
interface so-0/0/1.0;
interface lo0.1;
route-distinguisher 192.168.7.1:3;
provider-tunnel {
    rsvp-te {
        label-switched-path-template {
            default-template;
        }
    }
}
vrf-import green-red-blue-import;
vrf-target target:65000:3;
routing-options {
    auto-export;
}
protocols {
    bgp {
        group PE-CE {
            export BGP-export;
            neighbor 10.0.79.2 {
                peer-as 65003;
            }
        }
    }
    pim {
        rp {
            local {
                address 10.3.33.3;
            }
        }
        interface so-0/0/1.0 {
            mode sparse;
        }
        interface lo0.1 {
            mode sparse;
        }
    }
}
mvpn ;
}
```



```

    }
}

```

The relevant sample configuration for Router CE3 follows.

Router CE3

```

interfaces {
  so-0/0/1 {
    unit 0 {
      description "to PE3";
      family inet {
        address 10.0.79.2/30;
      }
    }
  }
  fe-0/1/0 {
    unit 0 {
      description "to H3";
      family inet {
        address 10.3.11.3/24;
      }
    }
  }
  lo0 {
    unit 0 {
      description "CE3 loopback";
      family inet {
        address 192.168.9.1/32 {
          primary;
        }
        address 127.0.0.1/32;
      }
    }
  }
}
routing-options {
  router-id 192.168.9.1;
  autonomous-system 65003;
  forwarding-table {
    export load-balance;
  }
}

```

```
}
protocols {
  bgp {
    group PE-CE {
      export BGP-export;
      neighbor 10.0.79.1 {
        peer-as 65000;
      }
    }
  }
  pim {
    rp {
      static {
        address 10.3.33.3;
      }
    }
    interface so-0/0/1.0 {
      mode sparse;
    }
    interface fe-0/1/0.0 {
      mode sparse;
    }
  }
}
policy-options {
  policy-statement BGP-export {
    term t1 {
      from protocol direct;
      then accept;
    }
    term t2 {
      from protocol static;
      then accept;
    }
  }
  policy-statement load-balance {
    then {
      load-balance per-packet;
    }
  }
}
```

RELATED DOCUMENTATION

[Configuring Multiprotocol BGP Multicast VPNs | 749](#)

[Multiprotocol BGP MVPNs Overview | 741](#)

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

In multiprotocol BGP (MBGP) 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.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

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.

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs

IN THIS SECTION

- [Requirements | 914](#)
- [Overview | 914](#)
- [Configuration | 915](#)
- [Verification | 925](#)

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

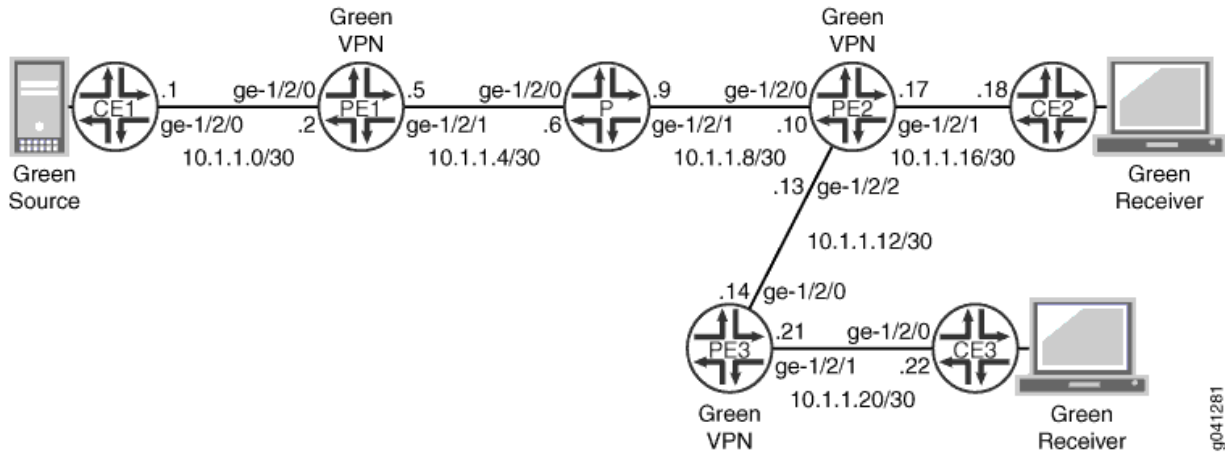
The routing device that has redundant VT interfaces configured must be running Junos OS Release 12.3 or later.

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 113 on page 915](#) shows the topology used in this example.

Figure 113: Multiple VT Interfaces in MBGP MVPN Topology



The following example shows the configuration for the customer edge (CE), provider (P), and provider edge (PE) devices in [Figure 113 on page 915](#). The section "[Step-by-Step Procedure](#)" on [page 920](#) describes the steps on Device PE2.

Configuration

IN THIS SECTION

- [Procedure | 915](#)

Procedure

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 10.1.1.1/30
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

```

Device CE2

```

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

```

Device CE3

```

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

```

Device P

```

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

```

Device PE1

```

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
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

IN THIS SECTION

- [Checking the LSP Route | 926](#)

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
                receive table vpn-1.inet.0, Pop
299872          *[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.

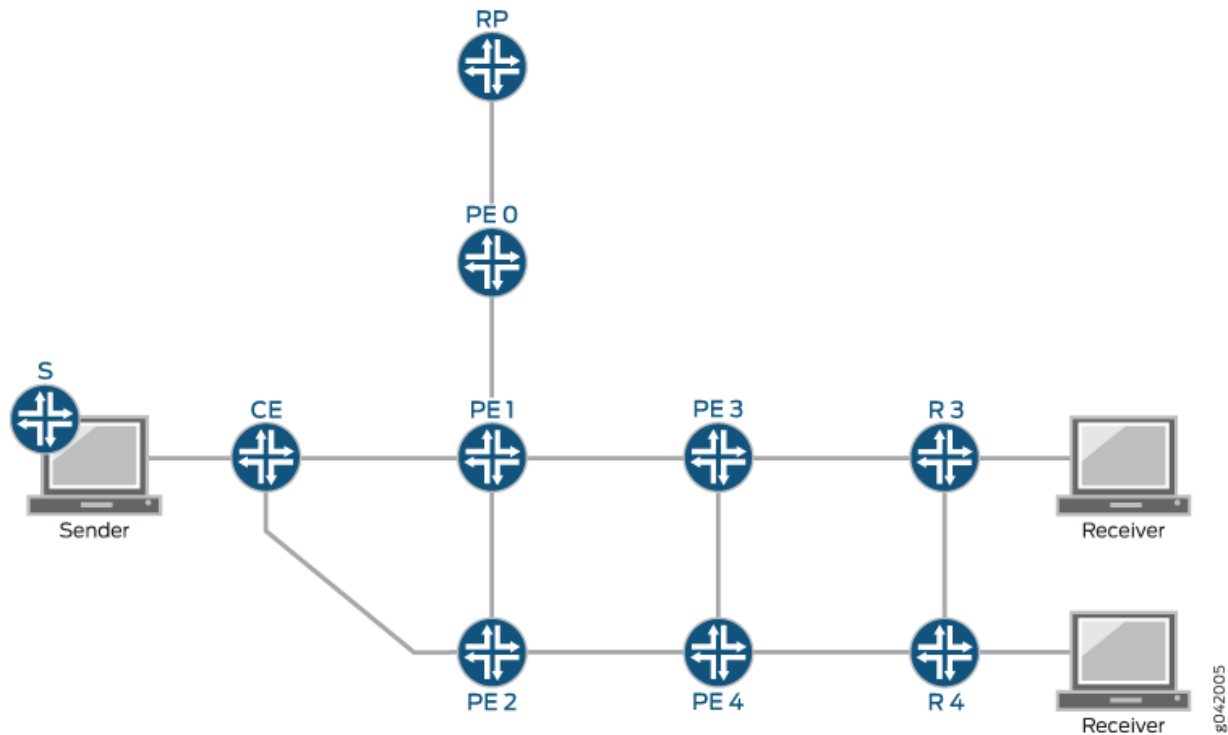
Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

IN THIS SECTION

- [Multiple Active and Backup Paths in RPF List | 932](#)

In a BGP multicast VPN (MVPN) (also called a multiprotocol BGP next-generation multicast VPN), sender-based reverse-path forwarding (RPF) helps to prevent multiple provider edge (PE) routers from sending traffic into the core, thus preventing duplicate traffic being sent to a customer. In the following diagram, sender-based RPF configured on egress Device PE3 and Device PE4 prevents duplicate traffic from being sent to the customers.

Figure 114: Sender-Based RPF



Sender-based RPF is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to `network-services enhanced-ip` mode.

Sender-based RPF (and hot-root standby) are supported only for MPLS BGP MVPNs with RSVP point-to-multipoint provider tunnels. Both SPT-only and SPT-RPT MVPN modes are supported.

Sender-based RPF does not work when point-to-multipoint provider tunnels are used with label-switched interfaces (LSI). Junos OS only allocates a single LSI label for each VRF, and uses this label for all point-to-multipoint tunnels. Therefore, the label that the egress receives does not indicate the sending PE router. LSI labels currently cannot scale to create a unique label for each point-to-multipoint tunnel. As such, virtual tunnel interfaces (vt) must be used for sender-based RPF functionality with point-to-multipoint provider tunnels.

Optionally, LSI interfaces can continue to be used for unicast purposes, and virtual tunnel interfaces can be configured to be used for multicast only.

Overview

In general, it is important to avoid (or recover from) having multiple PE routers send duplicate traffic into the core because this can result in duplicate traffic being sent to the customer. The sender-based RPF has a use case that is limited to BGP MVPNs. The use-case scope is limited for the following reasons:

- A traditional RPF check for native PIM is based on the incoming interface. This RPF check prevents loops but does not prevent multiple forwarders on a LAN. The traditional RPF has been used because current multicast protocols either avoid duplicates on a LAN or have data-driven events to resolve the duplicates once they are detected.
- In PIM sparse mode, duplicates can occur on a LAN in normal protocol operation. The protocol has a data-driven mechanism (PIM assert messages) to detect duplication when it happens and resolve it.
- In PIM bidirectional mode, a designated forwarder (DF) election is performed on all LANs to avoid duplication.
- Draft Rosen MVPNs use the PIM assert mechanism because with Draft Rosen MVPNs the core network is analogous to a LAN.

Sender-based RPF is a solution to be used in conjunction with BGP MVPNs because BGP MVPNs use an alternative to data-driven-event solutions and bidirectional mode DF election. This is so, because, for one thing, the core network is not exactly a LAN. In an MVPN scenario, it is possible to determine which PE router has sent the traffic. Junos OS uses this information to only forward the traffic if it is sent from the correct PE router. With sender-based RPF, the RPF check is enhanced to check whether data arrived on the correct incoming virtual tunnel (vt-) interface and that the data was sent from the correct upstream PE router.

More specifically, the data must arrive with the correct MPLS label in the outer header used to encapsulate data through the core. The label identifies the tunnel and, if the tunnel is point-to-multipoint, the upstream PE router.

Sender-based RPF is not a replacement for single-forwarder election, but is a complementary feature. Configuring a higher primary loopback address (or router ID) on one PE device (PE1) than on another (PE2) ensures that PE1 is the single-forwarder election winner. The `unicast-umh-election` statement causes the unicast route preference to determine the single-forwarder election. If single-forwarder election is not used or if it is not sufficient to prevent duplicates in the core, sender-based RPF is recommended.

For RSVP point-to-multipoint provider tunnels, the transport label identifies the sending PE router because it is a requirement that penultimate hop popping (PHP) is disabled when using point-to-multipoint provider tunnels with MVPNs. PHP is disabled by default when you configure the MVPN protocol in a routing instance. The label identifies the tunnel, and (because the RSVP-TE tunnel is point-to-multipoint) the sending PE router.

The sender-based RPF mechanism is described in RFC 6513, *Multicast in MPLS/BGP IP VPNs* in section 9.1.1.



NOTE: The hot-root standby technique described in Internet draft draft-morin-l3vpn-mvpn-fast-failover-05 *Multicast VPN fast upstream failover* is an egress PE router functionality in which the egress PE router sends source-tree c-multicast join message to

both a primary and a backup upstream PE router. This allows multiple copies of the traffic to flow through the provider core to the egress PE router. Sender-based RPF and hot-root standby can be used together to support *live-live* BGP MVPN traffic. This is a multicast-over-MPLS scheme for carrying mission-critical professional broadcast TV and IPTV traffic. A key requirement for many of these deployments is to have full redundancy of network equipment, including the ingress and egress PE routers. In some cases, a live-live approach is required, meaning that two duplicate traffic flows are sent across the network following diverse paths. When this technique is combined with sender-based forwarding, the two live flows of traffic are received at the egress PE router, and the egress PE router forwards a single stream to the customer network. Any failure in the network can be repaired locally at the egress PE router. For more information about hot-root standby, see [hot-root-standby](#).

Sender-based RPF prevents duplicates from being sent to the customer even if there is duplication in the provider network. Duplication could exist in the provider because of a hot-root standby configuration or if the single-forwarder election is not sufficient to prevent duplicates. Single-forwarder election is used to prevent duplicates to the core network, while sender-based RPF prevents duplicates to the customer even if there are duplicates in the core. There are cases in which single-forwarder election cannot prevent duplicate traffic from arriving at the egress PE router. One example of this (outlined in section 9.3.1 of RFC 6513) is when PIM sparse mode is configured in the customer network and the MVPN is in RPT-SPT mode with an I-PMSI.

Determining the Upstream PE Router

After Junos OS chooses the ingress PE router, the sender-based RPF decision determines whether the correct ingress PE router is selected. As described in RFC 6513, section 9.1.1, an egress PE router, PE1, chooses a specific upstream PE router, for given (C-S,C-G). When PE1 receives a (C-S,C-G) packet from a PMSI, it might be able to identify the PE router that transmitted the packet onto the PMSI. If that transmitter is other than the PE router selected by PE1 as the upstream PE router, PE1 can drop the packet. This means that the PE router detects a duplicate, but the duplicate is not forwarded.

When an egress PE router generates a type 7 C-multicast route, it uses the VRF route import extended community carried in the VPN-IP route toward the source to construct the route target carried by the C-multicast route. This route target results in the C-multicast route being sent to the upstream PE router, and being imported into the correct VRF on the upstream PE router. The egress PE router programs the forwarding entry to only accept traffic from this PE router, and only on a particular tunnel rooted at that PE router.

When an egress PE router generates a type 6 C-multicast route, it uses the VRF route import extended community carried in the VPN-IP route toward the rendezvous point (RP) to construct the route target carried by the C-multicast route.

This route target results in the C-multicast route being sent to the upstream PE router and being imported into the correct VRF on the upstream PE router. The egress PE router programs the forwarding entry to accept traffic from this PE router only, and only on a particular tunnel rooted at that PE router. However, if some other PE routers have switched to SPT mode for (C-S, C-G) and have sent source active (SA) autodiscovery (A-D) routes (type 5 routes), and if the egress PE router only has (C-*, C-G) state, the upstream PE router for (C-S, C-G) is not the PE router toward the RP to which it sent a type 6 route, but the PE router that originates a SA A-D route for (C-S, C-G). The traffic for (C-S, C-G) might be carried over a I-PMSI or S-PMSI, depending on how it was advertised by the upstream PE router.

Additionally, when an egress PE router has only the (C-*, C-G) state and does not have the (C-S, C-G) state, the egress PE router might be receiving (C-S, C-G) type 5 SA routes from multiple PE routers, and chooses the best one, as follows: For every received (C-S, C-G) SA route, the egress PE router finds in its upstream multicast hop (UMH) route-candidate set for C-S a route with the same route distinguisher (RD). Among all such found routes the PE router selects the UMH route (based on the UMH selection). The best (C-S, C-G) SA route is the one whose RD is the same as of the selected UMH route.

When an egress PE router has only the (C-*, C-G) state and does not have the (C-S, C-G) state, and if later the egress PE router creates the (C-S, C-G) state (for example, as a result of receiving a PIM join (C-S, C-G) message from one of its customer edge [CE] neighbors), the upstream PE router for that (C-S, C-G) is not necessarily going to be the same PE router that originated the already-selected best SA A-D route for (C-S, C-G). It is possible to have a situation in which the PE router that originated the best SA A-D route for (C-S, C-G) carries the (C-S, C-G) over an I-PMSI, while some other PE router, that is also connected to the site that contains C-S, carries (C-S,C-G) over an S-PMSI. In this case, the downstream PE router would not join the S-PMSI, but continue to receive (C-S, C-G) over the I-PMSI, because the UMH route for C-S is the one that has been advertised by the PE router that carries (C-S, C-G) over the I-PMSI. This is expected behavior.

The egress PE router determines the sender of a (C-S, C-G) type 5 SA A-D route by finding in its UMH route-candidate set for C-S a route whose RD is the same as in the SA A-D route. The VRF route import extended community of the found route contains the IP address of the sender of the SA A-D route.

Multiple Active and Backup Paths in RPF List

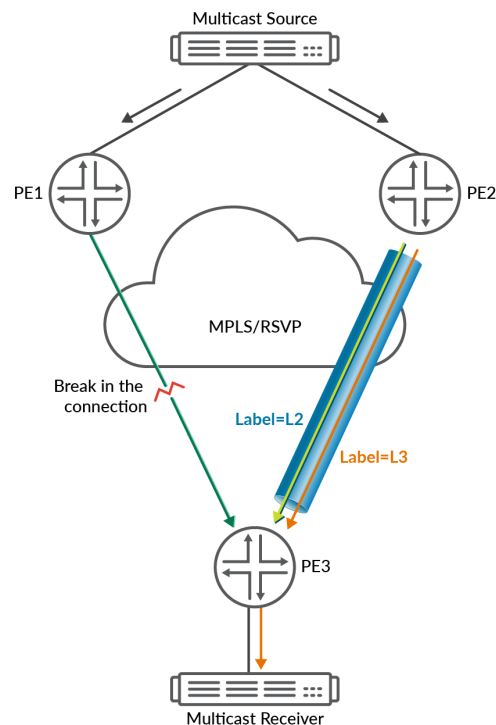
During a Make Before Break (MBB) event, Junos OS assigns multiple equal weight labels to a Label Switched Path (LSP) in an MVPN provider tunnel. The egress device accepts traffic only from the active next-hop i.e, the next-hop that was installed first. The next-hop installed thereafter is treated as the discard next-hop. As long as the traffic flows through the label installed first, there is no traffic loss. When the traffic from the ingress PE flows through the label installed next (treated as discard at the egress PE) there is a transient loss of traffic until the MBB event is completed.

Starting in Junos OS Evolved 23.4R1, a `Session Id` is created, based on the name of the provider tunnel. Sessions are grouped under this `Session Id` for a unicast next-hop. With this, different labels in the same LSP will be assigned the same `Session Id`. Junos OS uses this `Session Id` to accept and forward traffic from any of the labels with a matching `Session ID`.

In the figure below, PE3 is configured with MVPN Hot Root Standby (HRS) thereby fetching multicast traffic from both primary ingress device PE1 and secondary ingress device PE2. An RSVP P2MP tunnel is established between PE1 and PE3, with the incoming label L1. A second RSVP P2MP tunnel exists between PE2 and PE3, with the incoming label L2. If PE1 is unreachable for any reason, the egress device PE3 starts fetching traffic from PE2. In the case an MBB event is triggered, PE2 signals for a new LSP path and PE3 allocates a new incoming LSP label L3. During this time, the RPF list in PE3 is programmed with two incoming labels. The ingress PE decides when to switch traffic from the old label to the new one. When the traffic switches to the new label, the old label is torn down. PE3 modifies its RPF NH to label L3, following which the traffic flow is restored.


With the grouping of both labels, L2 and L3 under one Session Id, switching between the two LSP labels becomes seamless and causes a minimal transition delay of sub-50 ms.

Figure 115: MBB event triggered during HRS



Similarly, for MVPN provider tunnels with a threshold for I-PMSI traffic rate, the traffic flows through the I-PMSI tunnel until the threshold is exceeded, in which case the traffic switches to the S-PMSI tunnel. During this switchover from the I-PMSI to S-PMSI tunnel, you may experience traffic loss due to a change in the next-hop from which the egress PE was receiving and forwarding traffic.

Figure 116: I-PMSI to S-PMSI switchover



Junos uses the Session Id to group I-PMSI and S-PMSI next-hops together, minimizing the transition delay to sub-50 ms.

Running the command `show multicast route extensive instance instance` will include the Session Id and Session Status if present.

```
user@router> show multicast route extensive instance instance1
```

```
Instance: vrf4 Family: INET
Group: 233.252.0.1
Source: 172.16.0.1/32
Upstream rpf interface list:
  vt-5/0/0.0 (P)
    Session Id: 0x38a7 Session Status: Up
    Min-rate: 3000 kbps Weight: 1
    Sender Id: Label 24
  vt-5/0/0.0 (B)
    Session Id: 0x38a8 Session Status: Up
```



```
Min-rate: 3000 kbps Weight: 65533
      Sender Id: Label 23
Downstream interface list:
      et-5/1/5.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 349 kBps, 1465 pps, 1552316 packets
RPF Next-hop ID: 5326
Next-hop ID: 1048585
Upstream protocol: MVPN
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
```

RELATED DOCUMENTATION

| [*hot-root-standby \(MBGP MVPN\)*](#)

RELATED DOCUMENTATION

| [Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels | 935](#)

| [*unicast-umh-election*](#)

Example: Configuring Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels

IN THIS SECTION

- [Requirements | 936](#)
- [Overview | 936](#)
- [Set Commands for All Devices in the Topology | 938](#)
- [Configuring Device PE2 | 943](#)
- [Verification | 952](#)

This example shows how to configure sender-based reverse-path forwarding (RPF) in a BGP multicast VPN (MVPN). Sender-based RPF helps to prevent multiple provider edge (PE) routers from sending traffic into the core, thus preventing duplicate traffic being sent to a customer.

Requirements

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

Sender-based RPF is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to [network-services](#) enhanced-ip mode.

Sender-based RPF is supported only for MPLS BGP MVPNs with RSVP-TE point-to-multipoint provider tunnels. Both SPT-only and SPT-RPT MVPN modes are supported.

Sender-based RPF does not work when point-to-multipoint provider tunnels are used with label-switched interfaces (LSI). Junos OS only allocates a single LSI label for each VRF, and uses this label for all point-to-multipoint tunnels. Therefore, the label that the egress receives does not indicate the sending PE router. LSI labels currently cannot scale to create a unique label for each point-to-multipoint tunnel. As such, virtual tunnel interfaces (vt) must be used for sender-based RPF functionality with point-to-multipoint provider tunnels.

This example requires Junos OS Release 14.2 or later on the PE router that has sender-based RPF enabled.

Overview

IN THIS SECTION

- [Topology | 937](#)

This example shows a single autonomous system (intra-AS scenario) in which one source sends multicast traffic (group 224.1.1.1) into the VPN (VRF instance vpn-1). Two receivers subscribe to the group. They are connected to Device CE2 and Device CE3, respectively. RSVP point-to-multipoint LSPs with inclusive provider tunnels are set up among the PE routers. PIM (C-PIM) is configured on the PE-CE links.

For MPLS, the signaling control protocol used here is LDP. Optionally, you can use RSVP to signal both point-to-point and point-to-multipoint tunnels.

OSPF is used for interior gateway protocol (IGP) connectivity, though IS-IS is also a supported option. If you use OSPF, you must enable OSPF traffic engineering.

For testing purposes, routers are used to simulate the source and the receivers. Device PE2 and Device PE3 are configured to statically join the 224.1.1.1 group by using the `set protocols igmp interface interface-name static group 224.1.1.1` command. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the CE devices attached to the receivers, to make them listen to the multicast group address, the example uses `set protocols sap listen 224.1.1.1`. A ping command is used to send multicast traffic into the BGP MBPN.

Sender-based RPF is enabled on Device PE2, as follows:

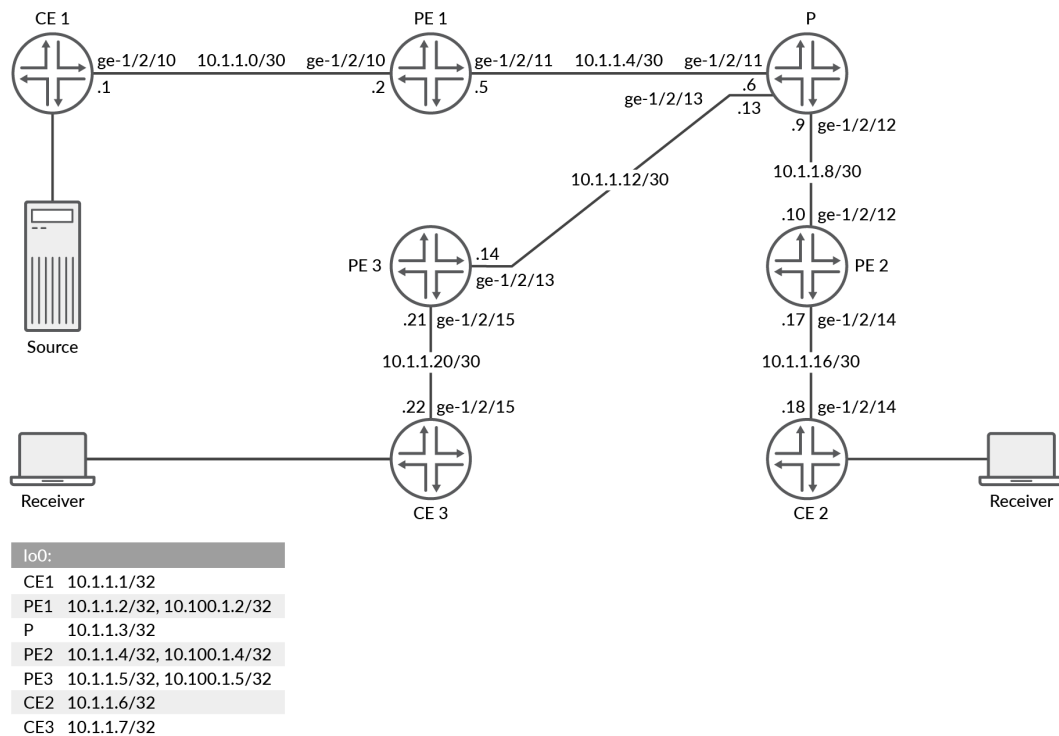
```
[routing-instances vpn-1 protocols mvpn]
user@PE2# set sender-based-rpf
```

You can optionally configure `hot-root-standby` with `sender-based-rpf`.

Topology

Figure 117 on page 937 shows the sample network.

Figure 117: Sender-Based RPF in a BGP MVPN



"Set Commands for All Devices in the Topology" on page 938 shows the configuration for all of the devices in Figure 117 on page 937.

The section "Configuring Device PE2" on page 943 describes the steps on Device PE2.

Set Commands for All Devices in the Topology

IN THIS SECTION

- [CLI Quick Configuration | 938](#)
- [Procedure | 943](#)

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-1/2/10 unit 0 family inet address 10.1.1.1/30
set interfaces ge-1/2/10 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/10.0
set protocols pim rp static address 10.100.1.2
set protocols pim interface all
set routing-options router-id 10.1.1.1
```

Device CE2

```
set interfaces ge-1/2/14 unit 0 family inet address 10.1.1.18/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/14.0
set protocols pim rp static address 10.100.1.2
```

```
set protocols pim interface all
set routing-options router-id 10.1.1.6
```

Device CE3

```
set interfaces ge-1/2/15 unit 0 family inet address 10.1.1.22/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.7/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/15.0
set protocols pim rp static address 10.100.1.2
set protocols pim interface all
set routing-options router-id 10.1.1.7
```

Device P

```
set interfaces ge-1/2/11 unit 0 family inet address 10.1.1.6/30
set interfaces ge-1/2/11 unit 0 family mpls
set interfaces ge-1/2/12 unit 0 family inet address 10.1.1.9/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/13 unit 0 family inet address 10.1.1.13/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.3/32
set protocols rsvp interface all
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls interface ge-1/2/11.0
set protocols mpls interface ge-1/2/12.0
set protocols mpls interface ge-1/2/13.0
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-1/2/11.0
set protocols ospf area 0.0.0.0 interface ge-1/2/12.0
set protocols ospf area 0.0.0.0 interface ge-1/2/13.0
set protocols ldp interface ge-1/2/11.0
set protocols ldp interface ge-1/2/12.0
set protocols ldp interface ge-1/2/13.0
set protocols ldp p2mp
set routing-options router-id 10.1.1.3
```

Device PE1

```

set interfaces ge-1/2/10 unit 0 family inet address 10.1.1.2/30
set interfaces ge-1/2/10 unit 0 family mpls
set interfaces ge-1/2/11 unit 0 family inet address 10.1.1.5/30
set interfaces ge-1/2/11 unit 0 family mpls
set interfaces vt-1/2/10 unit 2 family inet
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set interfaces lo0 unit 102 family inet address 10.100.1.2/32
set protocols rsvp interface ge-1/2/11.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path p2mp-template template
set protocols mpls label-switched-path p2mp-template p2mp
set protocols mpls interface ge-1/2/11.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.1.1.2
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.1.1.4
set protocols bgp group ibgp neighbor 10.1.1.5
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-1/2/11.0
set protocols ldp interface ge-1/2/11.0
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-1/2/10.0
set routing-instances vpn-1 interface vt-1/2/10.2
set routing-instances vpn-1 interface lo0.102
set routing-instances vpn-1 provider-tunnel rsvp-te label-switched-path-template p2mp-template
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0 rsvp-
te label-switched-path-template p2mp-template
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
threshold-rate 0
set routing-instances vpn-1 vrf-target target:100:10
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-1/2/10.0
set routing-instances vpn-1 protocols pim rp local address 10.100.1.2

```

```

set routing-instances vpn-1 protocols pim interface ge-1/2/10.0 mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode rpt-spt
set routing-options router-id 10.1.1.2
set routing-options route-distinguisher-id 10.1.1.2
set routing-options autonomous-system 65001

```

Device PE2

```

set interfaces ge-1/2/12 unit 0 family inet address 10.1.1.10/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/14 unit 0 family inet address 10.1.1.17/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces vt-1/2/10 unit 4 family inet
set interfaces lo0 unit 0 family inet address 10.1.1.4/32
set interfaces lo0 unit 104 family inet address 10.100.1.4/32
set protocols igmp interface ge-1/2/14.0 static group 224.1.1.1
set protocols rsvp interface ge-1/2/12.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path p2mp-template template
set protocols mpls label-switched-path p2mp-template p2mp
set protocols mpls interface ge-1/2/12.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.1.1.4
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.1.1.2
set protocols bgp group ibgp neighbor 10.1.1.5
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-1/2/12.0
set protocols ldp interface ge-1/2/12.0
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-1/2/10.4
set routing-instances vpn-1 interface ge-1/2/14.0
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 provider-tunnel rsvp-te label-switched-path-template p2mp-template
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0 rsvp-
te label-switched-path-template p2mp-template

```

```

set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
threshold-rate 0
set routing-instances vpn-1 vrf-target target:100:10
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-1/2/14.0
set routing-instances vpn-1 protocols pim rp static address 10.100.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/14.0 mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode rpt-spt
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 10.1.1.4
set routing-options route-distinguisher-id 10.1.1.4
set routing-options autonomous-system 65001

```

Device PE3

```

set interfaces ge-1/2/13 unit 0 family inet address 10.1.1.14/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 family inet address 10.1.1.21/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces vt-1/2/10 unit 5 family inet
set interfaces lo0 unit 0 family inet address 10.1.1.5/32
set interfaces lo0 unit 105 family inet address 10.100.1.5/32
set protocols igmp interface ge-1/2/15.0 static group 224.1.1.1
set protocols rsvp interface ge-1/2/13.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path p2mp-template template
set protocols mpls label-switched-path p2mp-template p2mp
set protocols mpls interface ge-1/2/13.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.1.1.5
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.1.1.2
set protocols bgp group ibgp neighbor 10.1.1.4
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-1/2/13.0
set protocols ldp interface ge-1/2/13.0
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-1/2/10.5
set routing-instances vpn-1 interface ge-1/2/15.0
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 provider-tunnel rsvp-te label-switched-path-template p2mp-template
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0 rsvp-
te label-switched-path-template p2mp-template
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
threshold-rate 0
set routing-instances vpn-1 vrf-target target:100:10
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-1/2/15.0
set routing-instances vpn-1 protocols pim rp static address 10.100.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/15.0 mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode rpt-spt
set routing-options router-id 10.1.1.5
set routing-options route-distinguisher-id 10.1.1.5
set routing-options autonomous-system 65001
```

Procedure

Step-by-Step Procedure

Configuring Device PE2

IN THIS SECTION

- [Procedure | 944](#)

Procedure

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 [Junos OS CLI User Guide](#).

To configure Device PE2:

1. Enable enhanced IP mode.

```
[edit chassis]
user@PE2# set network-services enhanced-ip
```

2. Configure the device interfaces.

```
[edit interfaces]
user@PE2# set ge-1/2/12 unit 0 family inet address 10.1.1.10/30
user@PE2# set ge-1/2/12 unit 0 family mpls
user@PE2# set ge-1/2/14 unit 0 family inet address 10.1.1.17/30
user@PE2# set ge-1/2/14 unit 0 family mpls
user@PE2# set vt-1/2/10 unit 4 family inet
user@PE2# set lo0 unit 0 family inet address 10.1.1.4/32
user@PE2# set lo0 unit 104 family inet address 10.100.1.4/32
```

3. Configure IGMP on the interface facing the customer edge.

```
[edit protocols igmp]
user@PE2# set interface ge-1/2/14.0
```

4. (Optional) Force the PE device to join the multicast group with a static configuration.

Normally, this would happen dynamically in a setup with real sources and receivers.

```
[edit protocols igmp]
user@PE2# set interface ge-1/2/14.0 static group 224.1.1.1
```

5. Configure RSVP on the interfaces facing the provider core.

```
[edit protocols rsvp]
user@PE2# set interface ge-1/2/0.10
```

6. Configure MPLS.

```
[edit protocols mpls]
user@PE2# set traffic-engineering bgp-igp-both-ribs
user@PE2# set label-switched-path p2mp-template template
user@PE2# set label-switched-path p2mp-template p2mp
user@PE2# set interface ge-1/2/12.0
```

7. Configure internal BGP (IBGP) among the PE routers.

```
[edit protocols bgp group ibgp]
user@PE2# set type internal
user@PE2# set local-address 10.1.1.4
user@PE2# set family inet unicast
user@PE2# set family inet-vpn any
user@PE2# set family inet-mvpn signaling
user@PE2# set neighbor 10.1.1.2
user@PE2# set neighbor 10.1.1.5
```

8. Configure an OSPF or IS-IS.

```
[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-1/2/12.0
```

9. (Optional) Configure LDP.

RSVP can be used instead for MPLS signaling.

```
[edit protocols bgp group ibgp]
user@PE2# set interface ge-1/2/12.0
user@PE2# set p2mp
```

10. Configure a routing policy to be used in the VPN.

The policy is used for exporting the BGP into the PE-CE IGP session.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE2# set from protocol bgp
user@PE2# set then accept
```

11. Configure the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set instance-type vrf
user@PE2# set interface vt-1/2/10.4
user@PE2# set interface ge-1/2/14.0
user@PE2# set interface lo0.104
```

12. Configure the provider tunnel.

```
[edit routing-instances vpn-1 provider-tunnel]
user@PE2# set rsvp-te label-switched-path-template p2mp-template
user@PE2# set selective group 225.0.1.0/24 source 0.0.0.0/0 rsvp-te label-switched-path-
template p2mp-template
user@PE2# set selective group 225.0.1.0/24 source 0.0.0.0/0 threshold-rate 0
```

13. Configure the VRF target.

In the context of unicast IPv4 routes, choosing `vrf-target` has two implications. First, every locally learned (in this case, direct and static) route at the VRF is exported to BGP with the specified route target (RT). Also, every received `inet-vpn` BGP route with that RT value is imported into the VRF `vpn-1`. This has the advantage of a simpler configuration, and the drawback of less flexibility in selecting and modifying the exported and imported routes. It also implies that the VPN is full mesh and all the PE routers get routes from each other, so complex configurations like hub-and-spoke or extranet are not feasible. If any of these features are required, it is necessary to use `vrf-import` and `vrf-export` instead.

```
[edit ]
user@PE2# set routing-instances vpn-1 vrf-target target:100:10
```

14. Configure the PE-CE OSPF session.

```
[edit routing-instances vpn-1 protocols ospf]
user@PE2# set export parent_vpn_routes
user@PE2# set area 0.0.0.0 interface lo0.104 passive
user@PE2# set area 0.0.0.0 interface ge-1/2/14.0
```

15. Configure the PE-CE PIM session.

```
[edit routing-instances vpn-1 protocols pim]
user@PE2# set rp static address 10.100.1.2
user@PE2# set interface ge-1/2/14.0 mode sparse
```

16. Enable the MVPN mode.

Both rpt-spt and spt-only are supported with sender-based RPF.

```
[edit routing-instances vpn-1 protocols mvpn]
user@PE2# set mvpn-mode rpt-spt
```

17. Enable sender-based RPF.

```
[edit routing-instances vpn-1 protocols mvpn]
user@PE2# set sender-based-rpf
```

18. Configure the router ID, the router distinguisher, and the AS number.

```
[edit routing-options]
user@PE2# set router-id 10.1.1.4
user@PE2# set route-distinguisher-id 10.1.1.4
user@PE2# set autonomous-system 65001
```

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 instructions in this example to correct the configuration.

```
user@PE2# show chassis
network-services enhanced-ip;
```

```
user@PE2# show interfaces
ge-1/2/12 {
  unit 0 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/14 {
  unit 0 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/2/10 {
  unit 5 {
    family inet;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.1.1.5/32;
    }
  }
  unit 105 {
    family inet {
      address 10.100.1.5/32;
    }
  }
}
```

```
}  
}
```

```
user@PE2# show protocols  
igmp {  
  interface ge-1/2/15.0 {  
    static {  
      group 224.1.1.1;  
    }  
  }  
}  
rsvp {  
  interface all;  
}  
mpls {  
  traffic-engineering bgp-igp-both-ribs;  
  label-switched-path p2mp-template {  
    template;  
    p2mp;  
  }  
  interface ge-1/2/13.0;  
}  
bgp {  
  group ibgp {  
    type internal;  
    local-address 10.1.1.5;  
    family inet {  
      unicast;  
    }  
    family inet-vpn {  
      any;  
    }  
    family inet-mvpn {  
      signaling;  
    }  
    neighbor 10.1.1.2;  
    neighbor 10.1.1.4;  
  }  
}  
ospf {  
  traffic-engineering;
```

```

area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
    interface ge-1/2/13.0;
}
}
ldp {
    interface ge-1/2/13.0;
    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-1/2/10.5;
    interface ge-1/2/15.0;
    interface lo0.105;
    provider-tunnel {
        rsvp-te {
            label-switched-path-template {
                p2mp-template;
            }
        }
        selective {
            group 225.0.1.0/24 {
                source 0.0.0.0/0 {
                    rsvp-te {
                        label-switched-path-template {
                            p2mp-template;
                        }
                    }
                }
                threshold-rate 0;
            }
        }
    }
}

```



```

    }
  }
}
vrf-target target:100:10;
protocols {
  ospf {
    export parent_vpn_routes;
    area 0.0.0.0 {
      interface lo0.105 {
        passive;
      }
      interface ge-1/2/15.0;
    }
  }
  pim {
    rp {
      static {
        address 100.1.1.2;
      }
    }
    interface ge-1/2/15.0 {
      mode sparse;
    }
  }
  mvpn {
    mvpn-mode {
      rpt-spt;
    }
    sender-based-rpf;
  }
}
}
}

```

```

user@PE2# show routing-options
router-id 10.1.1.5;
route-distinguisher-id 10.1.1.5;
autonomous-system 65001;

```

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

Verification

IN THIS SECTION

- [Verifying Sender-Based RPF | 952](#)
- [Checking the BGP Routes | 953](#)
- [Checking the PIM Joins on the Downstream CE Receiver Devices | 961](#)
- [Checking the PIM Joins on the PE Devices | 962](#)
- [Checking the Multicast Routes | 965](#)
- [Checking the MVPN C-Multicast Routes | 967](#)
- [Checking the Source PE | 968](#)

Confirm that the configuration is working properly.

Verifying Sender-Based RPF

Purpose

Make sure that sender-based RPF is enabled on Device PE2.

Action

```

user@PE2> show mvpn instance vpn-1

MVPN instance:
Legend for provider tunnel
S-   Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Family : INET

Instance : vpn-1
MVPN Mode : RPT-SPT
Sender-Based RPF: Enabled.
Hot Root Standby: Disabled. Reason: Not enabled by configuration.
Provider tunnel: I-P-tnl:RSVP-TE P2MP:10.1.1.4, 32647,10.1.1.4

```

```

Neighbor                               Inclusive Provider Tunnel
10.1.1.2                               RSVP-TE P2MP:10.1.1.2, 15282,10.1.1.2
10.1.1.5                               RSVP-TE P2MP:10.1.1.5, 8895,10.1.1.5
C-mcast IPv4 (S:G)                    Provider Tunnel                               St
0.0.0.0/0:224.1.1.1/32                RSVP-TE P2MP:10.1.1.2, 15282,10.1.1.2
0.0.0.0/0:224.2.127.254/32            RSVP-TE P2MP:10.1.1.2, 15282,10.1.1.2

```

MVPN instance:

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET6

Instance : vpn-1

MVPN Mode : RPT-SPT

Sender-Based RPF: Enabled.

Hot Root Standby: Disabled. Reason: Not enabled by configuration.

Provider tunnel: I-P-tnl:RSVP-TE P2MP:10.1.1.4, 32647,10.1.1.4

Checking the BGP Routes

Purpose

Make sure the expected BGP routes are being added to the routing tables on the PE devices.

Action

```

user@PE1> show route protocol bgp

inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

vpn-1.inet.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.6/32          *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.4
                    AS path: I, validation-state: unverified
                    > via ge-1/2/11.0, Push 299776, Push 299792(top)

```

```

10.1.1.7/32      *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.1.1.16/30    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.20/30    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.100.1.4/32   *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.100.1.5/32   *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)

```

vpn-1.inet.1: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```

10.1.1.4:32767:10.1.1.6/32
                 *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.1.1.16/30
                 *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.100.1.4/32
                 *[BGP/170] 1d 04:23:24, localpref 100, from .1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.5:32767:10.1.1.7/32
                 *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.1.1.20/30
                 *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)

```

```

1.1.1.5:32767:100.1.1.5/32
    *[BGP/170] 1d 04:23:23, localpref 100, from 1.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776, Push 299776(top)

bgp.mvpn.0: 5 destinations, 8 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.4:32767:10.1.1.4/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792

1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776

6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.1.1.1/240
    *[BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 1d 04:17:24, MED 0, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792

6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.2.127.254/240
    *[BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 1d 04:17:23, MED 0, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792

7:10.1.1.2:32767:1001:32:10.1.1.1:32:224.1.1.1/240
    *[BGP/170] 20:34:47, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 20:34:47, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792

vpn-1.mvpn.0: 7 destinations, 13 routes (7 active, 2 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.4:32767:10.1.1.4/240

```

```

*[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299792
1:10.1.1.5:32767:10.1.1.5/240
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299776
6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.1.1.1/240
[BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299776
[BGP/170] 1d 04:17:24, MED 0, localpref 100, from 10.1.1.4
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299792
6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.2.127.254/240
[BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299776
[BGP/170] 1d 04:17:23, MED 0, localpref 100, from 10.1.1.4
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299792
7:10.1.1.2:32767:1001:32:10.1.1.1:32:224.1.1.1/240
[BGP/170] 20:34:47, localpref 100, from 10.1.1.4
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299792
[BGP/170] 20:34:47, localpref 100, from 10.1.1.5
  AS path: I, validation-state: unverified
  > via ge-1/2/11.0, Push 299776

```

```
user@PE2> show route protocol bgp
```

```
inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)
```

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
vpn-1.inet.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.1.1.1/32      *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.2
                  AS path: I, validation-state: unverified
```

```

> via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.7/32 * [BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.5
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299776(top)
10.1.1.0/30 * [BGP/170] 1d 04:23:24, localpref 100, from 1.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.20/30 * [BGP/170] 1d 04:23:20, localpref 100, from 1.1.1.5
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299776(top)
10.100.1.2/32 * [BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299808(top)
10.100.1.5/32 * [BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299776(top)

vpn-1.inet.1: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2:32767:10.1.1.1/32
* [BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.1.1.0/30
* [BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.100.1.2/32
* [BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.5:32767:10.1.1.7/32
* [BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.5
AS path: I, validation-state: unverified
> via ge-1/2/12.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.1.1.20/30
* [BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
AS path: I, validation-state: unverified

```

```

    > via ge-1/2/12.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.100.1.5/32
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/12.0, Push 299776, Push 299776(top)

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

10:1.1.1.2:32767:10.1.1.2/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
    AS path: I, validation-state: unverified
    > via ge-1/2/12.0, Push 299808
1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/12.0, Push 299776
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
    *[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
    AS path: I, validation-state: unverified
    > via ge-1/2/12.0, Push 299808

vpn-1.mvpn.0: 7 destinations, 9 routes (7 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.2:32767:10.1.1.2/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
    AS path: I, validation-state: unverified
    > via ge-1/2/12.0, Push 299808
1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/12.0, Push 299776
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
    *[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
    AS path: I, validation-state: unverified

```



```
> via ge-1/2/12.0, Push 299808
```

```
user@PE3> show route protocol bgp
```

```
inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)
```

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
vpn-1.inet.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.1.1.1/32      *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299808(top)
```

```
10.1.1.6/32      *[BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299792(top)
```

```
10.1.1.0/30      *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299808(top)
```

```
10.1.1.16/30     *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299792(top)
```

```
10.100.1.2/32    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299808(top)
```

```
10.100.1.4/32    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299792(top)
```

```
vpn-1.inet.1: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
```

```
bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.1.1.2:32767:10.1.1.1/32
```

```
      *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.2
      AS path: I, validation-state: unverified
```

```

> via ge-1/2/13.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.1.1.0/30
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299808(top)
10.1.1.2:32767:100.1.1.2/32
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299808(top)
10.1.1.4:32767:1.01.1.6/32
*[BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.1.1.16/30
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.100.1.4/32
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299792(top)

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

1:10.1.1.2:32767:10.1.1.2/240
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299808
1:10.1.1.4:32767:10.1.1.4/240
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299792
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
*[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299808

vpn-1.mvpn.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.2:32767:10.1.1.2/240
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2

```

```

AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299808
1:10.1.1.4:32767:10.1.1.4/240
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299792
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
*[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299808

```

Checking the PIM Joins on the Downstream CE Receiver Devices

Purpose

Make sure that the expected join messages are being sent.

Action

```

user@CE2> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
Source: *
RP: 10.100.1.2
Flags: sparse,rptree,wildcard
Upstream interface: ge-1/2/14.0

Group: 224.2.127.254
Source: *
RP: 10.100.1.2
Flags: sparse,rptree,wildcard
Upstream interface: ge-1/2/14.0

Instance: PIM.master Family: INET6

```

```
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
-----
```

```
user@CE3> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-1/2/15.0

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-1/2/15.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
-----
```

Meaning

Both Device CE2 and Device CE3 send C-Join packets upstream to their neighboring PE routers, their unicast next-hop to reach the C-Source.

Checking the PIM Joins on the PE Devices

Purpose

Make sure that the expected join messages are being sent.

Action

```
user@PE1> show pim join instance vpn-1
Instance: PIM.vpn-1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
```

```
Group: 224.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream interface: ge-1/2/10.0
```

```
Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
```

```
user@PE2> show pim join instance vpn-1
Instance: PIM.vpn-1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP
```

```
Group: 224.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream protocol: BGP
  Upstream interface: Through BGP
```

```
Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
```

Upstream interface: Through BGP

```
user@PE3> show pim join instance vpn-1
Instance: PIM.vpn-1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP

Group: 224.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream protocol: BGP
  Upstream interface: Through BGP

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP
```

Meaning

Both Device CE2 and Device CE3 send C-Join packets upstream to their neighboring PE routers, their unicast next-hop to reach the C-Source.

The C-Join state points to BGP as the upstream interface. Actually, there is no PIM neighbor relationship between the PEs. The downstream PE converts the C-PIM (C-S, C-G) state into a Type 7 source-tree join BGP route, and sends it to the upstream PE router toward the C-Source.

Checking the Multicast Routes

Purpose

Make sure that the C-Multicast flow is integrated in MVPN vpn-1 and sent by Device PE1 into the provider tunnel.

Action

```
user@PE1> show multicast route instance vpn-1
```

```
Instance: vpn-1 Family: INET
```

```
Group: 224.1.1.1/32
```

```
Source: *
```

```
Upstream interface: local
```

```
Downstream interface list:
```

```
ge-1/2/11.0
```

```
Group: 224.1.1.1
```

```
Source: 10.1.1.1/32
```

```
Upstream interface: ge-1/2/10.0
```

```
Downstream interface list:
```

```
ge-1/2/11.0
```

```
Group: 224.2.127.254/32
```

```
Source: *
```

```
Upstream interface: local
```

```
Downstream interface list:
```

```
ge-1/2/11.0
```

```
user@PE2> show multicast route instance vpn-1
```

```
Instance: vpn-1 Family: INET
```

```
Group: 224.1.1.1/32
```

```
Source: *
```

```
Upstream rpf interface list:
```

```
vt-1/2/10.4 (P)
```

```
Sender Id: Label 299840
```

```
Downstream interface list:
```

```
ge-1/2/14.0
```

```

Group: 224.1.1.1
  Source: 10.1.1.1/32
  Upstream rpf interface list:
    vt-1/2/10.4 (P)
      Sender Id: Label 299840

```

```

Group: 224.2.127.254/32
  Source: *
  Upstream rpf interface list:
    vt-1/2/10.4 (P)
      Sender Id: Label 299840
  Downstream interface list:
    ge-1/2/14.0

```

```
user@PE3> show multicast route instance vpn-1
```

```
Instance: vpn-1 Family: INET
```

```

Group: 224.1.1.1/32
  Source: *
  Upstream interface: vt-1/2/10.5
  Downstream interface list:
    ge-1/2/15.0

```

```

Group: 224.1.1.1
  Source: 10.1.1.1/32
  Upstream interface: vt-1/2/10.5

```

```

Group: 224.2.127.254/32
  Source: *
  Upstream interface: vt-1/2/10.5
  Downstream interface list:
    ge-1/2/15.0

```

Meaning

The output shows that, unlike the other PE devices, Device PE2 is using sender-based RPF. The output on Device PE2 includes the upstream RPF sender. The Sender Id field is only shown when sender-based RPF is enabled.

Checking the MVPN C-Multicast Routes

Purpose

Check the MVPN C-multicast route information,

Action

```
user@PE1> show mvpn c-multicast instance-name vpn-1
```

MVPN instance:

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET

Instance : vpn-1

MVPN Mode : RPT-SPT

C-mcast IPv4 (S:G)	Provider Tunnel	St
0.0.0.0/0:224.1.1.1/32	RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2	RM
10.1.1.1/32:224.1.1.1/32	RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2	RM
0.0.0.0/0:224.2.127.254/32	RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2	RM

...

```
user@PE2> show mvpn c-multicast instance-name vpn-1
```

MVPN instance:

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET

Instance : vpn-1

MVPN Mode : RPT-SPT

C-mcast IPv4 (S:G)	Provider Tunnel	St
--------------------	-----------------	----

```

0.0.0.0/0:224.1.1.1/32      RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2
10.1.1.1/32:224.1.1.1/32  RSVP-TE P2MP:1.1.1.2, 33314,10.1.1.2
0.0.0.0/0:224.2.127.254/32  RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2
...

```

```

user@PE3> show mvpn c-multicast instance-name vpn-1

MVPN instance:
Legend for provider tunnel
S-   Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Family : INET

Instance : vpn-1
MVPN Mode : RPT-SPT
C-mcast IPv4 (S:G)      Provider Tunnel      St
0.0.0.0/0:224.1.1.1/32  RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2
10.1.1.1/32:224.1.1.1/32  RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2
0.0.0.0/0:224.2.127.254/32  RSVP-TE P2MP:10.1.1.2, 33314,10.1.1.2
...

```

Meaning

The output shows the provider tunnel and label information.

Checking the Source PE

Purpose

Check the details of the source PE,

Action

```

user@PE1> show mvpn c-multicast source-pe

```

```

Instance : vpn-1
MVPN Mode : RPT-SPT
Family : INET
  C-Multicast route address :0.0.0.0/0:224.1.1.1/32
    MVPN Source-PE1:
      extended-community: no-advertise target:10.1.1.2:72
      Route Distinguisher: 10.1.1.2:32767
      Autonomous system number: 65001
      Interface: lo0.102 Index: -1610691384
    PIM Source-PE1:
      extended-community: target:10.1.1.2:72
      Route Distinguisher: 10.1.1.2:32767
      Autonomous system number: 65001
      Interface: lo0.102 Index: -1610691384
  C-Multicast route address :10.1.1.1/32:224.1.1.1/32
    MVPN Source-PE1:
      extended-community: no-advertise target:10.1.1.2:72
      Route Distinguisher: 10.1.1.2:32767
      Autonomous system number: 65001
      Interface: ge-1/2/10.0 Index: -1610691384
    PIM Source-PE1:
      extended-community: target:10.1.1.2:72
      Route Distinguisher: 10.1.1.2:32767
      Autonomous system number: 65001
      Interface: ge-1/2/10.0 Index: -1610691384
  C-Multicast route address :0.0.0.0/0:224.2.127.254/32
    MVPN Source-PE1:
      extended-community: no-advertise target:10.1.1.2:72
      Route Distinguisher: 10.1.1.2:32767
      Autonomous system number: 65001
      Interface: lo0.102 Index: -1610691384
    PIM Source-PE1:
      extended-community: target:10.1.1.2:72
      Route Distinguisher: 10.1.1.2:32767
      Autonomous system number: 65001
      Interface: lo0.102 Index: -1610691384

```

```
user@PE2> show mvpn c-multicast source-pe
```

```

Instance : vpn-1
MVPN Mode : RPT-SPT
Family : INET

```

```

C-Multicast route address :0.0.0.0/0:224.1.1.1/32
  MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
  PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
C-Multicast route address :10.1.1.1/32:224.1.1.1/32
  MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
  PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
C-Multicast route address :0.0.0.0/0:224.2.127.254/32
  MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
  PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)

```

```
user@PE3> show mvpn c-multicast source-pe
```

```

Instance : vpn-1
MVPN Mode : RPT-SPT
Family : INET
C-Multicast route address :0.0.0.0/0:224.1.1.1/32

```

```

MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
C-Multicast route address :10.1.1.1/32:224.1.1.1/32
MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
C-Multicast route address :0.0.0.0/0:224.2.127.254/32
MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
...

```

Meaning

The output shows the provider tunnel and label information.

RELATED DOCUMENTATION

[Understanding Sender-Based RPF in a BGP MVPN with RSVP-TE Point-to-Multipoint Provider Tunnels | 928](#)

unicast-umh-election

Example: Configuring Sender-Based RPF in a BGP MVPN with MLDP Point-to-Multipoint Provider Tunnels

IN THIS SECTION

- [Requirements | 972](#)
- [Overview | 973](#)
- [Set Commands for All Devices in the Topology | 974](#)
- [Configuring Device PE2 | 980](#)
- [Verification | 988](#)

This example shows how to configure sender-based reverse-path forwarding (RPF) in a BGP multicast VPN (MVPN). Sender-based RPF helps to prevent multiple provider edge (PE) routers from sending traffic into the core, thus preventing duplicate traffic being sent to a customer.

Requirements

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

Sender-based RPF is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to [network-services](#) enhanced-ip mode.

Sender-based RPF is supported only for MPLS BGP MVPNs with RSVP-TE point-to-multipoint provider tunnels. Both SPT-only and SPT-RPT MVPN modes are supported.

Sender-based RPF does not work when point-to-multipoint provider tunnels are used with label-switched interfaces (LSI). Junos OS only allocates a single LSI label for each VRF, and uses this label for all point-to-multipoint tunnels. Therefore, the label that the egress receives does not indicate the sending PE router. LSI labels currently cannot scale to create a unique label for each point-to-multipoint tunnel. As such, virtual tunnel interfaces (vt) must be used for sender-based RPF functionality with point-to-multipoint provider tunnels.

This example requires Junos OS Release 21.1R1 or later on the PE router that has sender-based RPF enabled.

Overview

IN THIS SECTION

- [Topology | 973](#)

This example shows a single autonomous system (intra-AS scenario) in which one source sends multicast traffic (group 224.1.1.1) into the VPN (VRF instance `vpn-1`). Two receivers subscribe to the group. They are connected to Device CE2 and Device CE3, respectively. MLDP point-to-multipoint LSPs with inclusive provider tunnels are set up among the PE routers. PIM (C-PIM) is configured on the PE-CE links.

For MPLS, the signaling control protocol used here is LDP. Optionally, you can use RSVP to signal both point-to-point and point-to-point tunnels.

OSPF is used for interior gateway protocol (IGP) connectivity, though IS-IS is also a supported option. If you use OSPF, you must enable OSPF traffic engineering.

For testing purposes, routers are used to simulate the source and the receivers. Device PE2 and Device PE3 are configured to statically join the 224.1.1.1 group by using the `set protocols igmp interface interface-name static group 224.1.1.1` command. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the CE devices attached to the receivers, to make them listen to the multicast group address, the example uses `set protocols sap listen 224.1.1.1`. A `ping` command is used to send multicast traffic into the BGP MBPN.

Sender-based RPF is enabled on Device PE2, as follows:

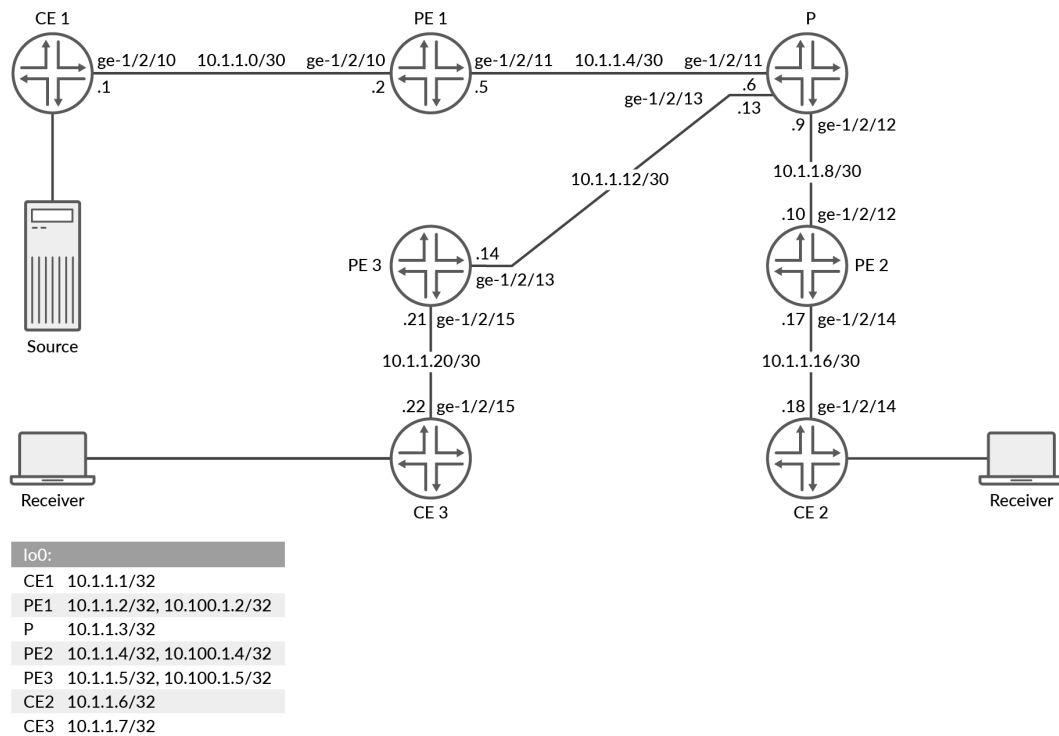
```
[routing-instances vpn-1 protocols mvpn]
user@PE2# set sender-based-rpf
```

You can optionally configure `hot-root-standby` with `sender-based-rpf`.

Topology

[Figure 118 on page 974](#) shows the sample network.

Figure 118: Sender-Based RPF in a BGP MVPN



"Set Commands for All Devices in the Topology" on page 974 shows the configuration for all of the devices in Figure 118 on page 974.

The section "Configuring Device PE2" on page 980 describes the steps on Device PE2.

Set Commands for All Devices in the Topology

IN THIS SECTION

- CLI Quick Configuration | 974
- Procedure | 980

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-1/2/10 unit 0 family inet address 10.1.1.1/30
set interfaces ge-1/2/10 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/10.0
set protocols pim rp static address 10.100.1.2
set protocols pim interface all
set routing-options router-id 10.1.1.1

```

Device CE2

```

set interfaces ge-1/2/14 unit 0 family inet address 10.1.1.18/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/14.0
set protocols pim rp static address 10.100.1.2
set protocols pim interface all
set routing-options router-id 10.1.1.6

```

Device CE3

```

set interfaces ge-1/2/15 unit 0 family inet address 10.1.1.22/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.7/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/15.0
set protocols pim rp static address 10.100.1.2
set protocols pim interface all
set routing-options router-id 10.1.1.7

```

Device P

```

set interfaces ge-1/2/11 unit 0 family inet address 10.1.1.6/30
set interfaces ge-1/2/11 unit 0 family mpls
set interfaces ge-1/2/12 unit 0 family inet address 10.1.1.9/30

```

```

set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/13 unit 0 family inet address 10.1.1.13/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.3/32
set protocols rsvp interface all
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls interface ge-1/2/11.0
set protocols mpls interface ge-1/2/12.0
set protocols mpls interface ge-1/2/13.0
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-1/2/11.0
set protocols ospf area 0.0.0.0 interface ge-1/2/12.0
set protocols ospf area 0.0.0.0 interface ge-1/2/13.0
set protocols ldp interface ge-1/2/11.0
set protocols ldp interface ge-1/2/12.0
set protocols ldp interface ge-1/2/13.0
set protocols ldp p2mp
set routing-options router-id 10.1.1.3

```

Device PE1

```

set interfaces ge-1/2/10 unit 0 family inet address 10.1.1.2/30
set interfaces ge-1/2/10 unit 0 family mpls
set interfaces ge-1/2/11 unit 0 family inet address 10.1.1.5/30
set interfaces ge-1/2/11 unit 0 family mpls
set interfaces vt-1/2/10 unit 2 family inet
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set interfaces lo0 unit 102 family inet address 10.100.1.2/32
set protocols rsvp interface ge-1/2/11.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path p2mp-template template
set protocols mpls label-switched-path p2mp-template p2mp
set protocols mpls interface ge-1/2/11.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.1.1.2
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.1.1.4
set protocols bgp group ibgp neighbor 10.1.1.5
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-1/2/11.0
set protocols ldp interface ge-1/2/11.0
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-1/2/10.0
set routing-instances vpn-1 interface vt-1/2/10.2
set routing-instances vpn-1 interface lo0.102
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0 ldp-
p2mp
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
threshold-rate 0
set routing-instances vpn-1 vrf-target target:100:10
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-1/2/10.0
set routing-instances vpn-1 protocols pim rp local address 10.100.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/10.0 mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode rpt-spt
set routing-options router-id 10.1.1.2
set routing-options route-distinguisher-id 10.1.1.2
set routing-options autonomous-system 65001

```

Device PE2

```

set interfaces ge-1/2/12 unit 0 family inet address 10.1.1.10/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/14 unit 0 family inet address 10.1.1.17/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces vt-1/2/10 unit 4 family inet
set interfaces lo0 unit 0 family inet address 10.1.1.4/32
set interfaces lo0 unit 104 family inet address 10.100.1.4/32
set protocols igmp interface ge-1/2/14.0 static group 224.1.1.1
set protocols rsvp interface ge-1/2/12.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path p2mp-template template
set protocols mpls label-switched-path p2mp-template p2mp
set protocols mpls interface ge-1/2/12.0
set protocols bgp group ibgp type internal

```

```

set protocols bgp group ibgp local-address 10.1.1.4
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.1.1.2
set protocols bgp group ibgp neighbor 10.1.1.5
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-1/2/12.0
set protocols ldp interface ge-1/2/12.0
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-1/2/10.4
set routing-instances vpn-1 interface ge-1/2/14.0
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0 ldp-
p2mp
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
threshold-rate 0
set routing-instances vpn-1 vrf-target target:100:10
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-1/2/14.0
set routing-instances vpn-1 protocols pim rp static address 10.100.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/14.0 mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode rpt-spt
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 10.1.1.4
set routing-options route-distinguisher-id 10.1.1.4
set routing-options autonomous-system 65001

```

Device PE3

```

set interfaces ge-1/2/13 unit 0 family inet address 10.1.1.14/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 family inet address 10.1.1.21/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces vt-1/2/10 unit 5 family inet

```

```
set interfaces lo0 unit 0 family inet address 10.1.1.5/32
set interfaces lo0 unit 105 family inet address 10.100.1.5/32
set protocols igmp interface ge-1/2/15.0 static group 224.1.1.1
set protocols rsvp interface ge-1/2/13.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path p2mp-template template
set protocols mpls label-switched-path p2mp-template p2mp
set protocols mpls interface ge-1/2/13.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.1.1.5
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 10.1.1.2
set protocols bgp group ibgp neighbor 10.1.1.4
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-1/2/13.0
set protocols ldp interface ge-1/2/13.0
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-1/2/10.5
set routing-instances vpn-1 interface ge-1/2/15.0
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
set routing-instances vpn-1 provider-tunnel selective group 225.0.1.0/24 source 0.0.0.0/0
threshold-rate 0
set routing-instances vpn-1 vrf-target target:100:10
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-1/2/15.0
set routing-instances vpn-1 protocols pim rp static address 10.100.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/15.0 mode sparse
set routing-instances vpn-1 protocols mvpn mvpn-mode rpt-spt
set routing-options router-id 10.1.1.5
set routing-options route-distinguisher-id 10.1.1.5
set routing-options autonomous-system 65001
```

Procedure

Step-by-Step Procedure

Configuring Device PE2

IN THIS SECTION

- Procedure | 980

Procedure

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 [Junos OS CLI User Guide](#).

To configure Device PE2:

1. Enable enhanced IP mode.

```
[edit chassis]
user@PE2# set network-services enhanced-ip
```

2. Configure the device interfaces.

```
[edit interfaces]
user@PE2# set ge-1/2/12 unit 0 family inet address 10.1.1.10/30
user@PE2# set ge-1/2/12 unit 0 family mpls
user@PE2# set ge-1/2/14 unit 0 family inet address 10.1.1.17/30
user@PE2# set ge-1/2/14 unit 0 family mpls
user@PE2# set vt-1/2/10 unit 4 family inet
user@PE2# set lo0 unit 0 family inet address 10.1.1.4/32
user@PE2# set lo0 unit 104 family inet address 10.100.1.4/32
```

3. Configure IGMP on the interface facing the customer edge.

```
[edit protocols igmp]
user@PE2# set interface ge-1/2/14.0
```

4. (Optional) Force the PE device to join the multicast group with a static configuration.

Normally, this would happen dynamically in a setup with real sources and receivers.

```
[edit protocols igmp]
user@PE2# set interface ge-1/2/14.0 static group 224.1.1.1
```

5. Configure RSVP on the interfaces facing the provider core.

```
[edit protocols rsvp]
user@PE2# set interface ge-1/2/0.10
```

6. Configure MPLS.

```
[edit protocols mpls]
user@PE2# set traffic-engineering bgp-igp-both-ribs
user@PE2# set label-switched-path p2mp-template template
user@PE2# set label-switched-path p2mp-template p2mp
user@PE2# set interface ge-1/2/12.0
```

7. Configure internal BGP (IBGP) among the PE routers.

```
[edit protocols bgp group ibgp]
user@PE2# set type internal
user@PE2# set local-address 10.1.1.4
user@PE2# set family inet unicast
user@PE2# set family inet-vpn any
user@PE2# set family inet-mvpn signaling
user@PE2# set neighbor 10.1.1.2
user@PE2# set neighbor 10.1.1.5
```

8. Configure an OSPF or IS-IS.

```
[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-1/2/12.0
```

9. (Optional) Configure LDP.

RSVP can be used instead for MPLS signaling.

```
[edit protocols bgp group ibgp]
user@PE2# set interface ge-1/2/12.0
user@PE2# set p2mp
```

10. Configure a routing policy to be used in the VPN.

The policy is used for exporting the BGP into the PE-CE IGP session.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE2# set from protocol bgp
user@PE2# set then accept
```

11. Configure the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set instance-type vrf
user@PE2# set interface vt-1/2/10.4
user@PE2# set interface ge-1/2/14.0
user@PE2# set interface lo0.104
```

12. Configure the provider tunnel.

```
[edit routing-instances vpn-1 provider-tunnel]
user@PE2# set ldp-p2mp
user@PE2# set selective group 225.0.1.0/24 source 0.0.0.0/0 ldp-p2mp
user@PE2# set selective group 225.0.1.0/24 source 0.0.0.0/0 threshold-rate 0
```

13. Configure the VRF target.

In the context of unicast IPv4 routes, choosing `vrf-target` has two implications. First, every locally learned (in this case, direct and static) route at the VRF is exported to BGP with the specified route target (RT). Also, every received `inet-vpn` BGP route with that RT value is imported into the VRF `vpn-1`. This has the advantage of a simpler configuration, and the drawback of less flexibility in selecting and modifying the exported and imported routes. It also implies that the VPN is full mesh and all the PE routers get routes from each other, so complex configurations like hub-and-spoke or extranet are not feasible. If any of these features are required, it is necessary to use `vrf-import` and `vrf-export` instead.

```
[edit ]
user@PE2# set routing-instances vpn-1 vrf-target target:100:10
```

14. Configure the PE-CE OSPF session.

```
[edit routing-instances vpn-1 protocols ospf]
user@PE2# set export parent_vpn_routes
user@PE2# set area 0.0.0.0 interface lo0.104 passive
user@PE2# set area 0.0.0.0 interface ge-1/2/14.0
```

15. Configure the PE-CE PIM session.

```
[edit routing-instances vpn-1 protocols pim]
user@PE2# set rp static address 10.100.1.2
user@PE2# set interface ge-1/2/14.0 mode sparse
```

16. Enable the MVPN mode.

Both `rpt-spt` and `spt-only` are supported with sender-based RPF.

```
[edit routing-instances vpn-1 protocols mvpn]
user@PE2# set mvpn-mode rpt-spt
```

17. Enable sender-based RPF.

```
[edit routing-instances vpn-1 protocols mvpn]
user@PE2# set sender-based-rpf
```

18. Configure the router ID, the router distinguisher, and the AS number.

```
[edit routing-options]
user@PE2# set router-id 10.1.1.4
user@PE2# set route-distinguisher-id 10.1.1.4
user@PE2# set autonomous-system 65001
```

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 instructions in this example to correct the configuration.

```
user@PE2# show chassis
network-services enhanced-ip;
```

```
user@PE2# show interfaces
ge-1/2/12 {
  unit 0 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/14 {
  unit 0 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/2/10 {
  unit 5 {
    family inet;
  }
}
```

```
lo0 {
  unit 0 {
    family inet {
      address 10.1.1.5/32;
    }
  }
  unit 105 {
    family inet {
      address 10.100.1.5/32;
    }
  }
}
```

```
user@PE2# show protocols
igmp {
  interface ge-1/2/15.0 {
    static {
      group 224.1.1.1;
    }
  }
}
rsvp {
  interface all;
}
mpls {
  traffic-engineering bgp-igp-both-ribs;
  label-switched-path p2mp-template {
    template;
    p2mp;
  }
  interface ge-1/2/13.0;
}
bgp {
  group ibgp {
    type internal;
    local-address 10.1.1.5;
    family inet {
      unicast;
    }
    family inet-vpn {
      any;
    }
  }
}
```

```

    }
    family inet-mvpn {
        signaling;
    }
    neighbor 10.1.1.2;
    neighbor 10.1.1.4;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface ge-1/2/13.0;
    }
}
}
ldp {
    interface ge-1/2/13.0;
    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-1/2/10.5;
    interface ge-1/2/15.0;
    interface lo0.105;
    provider-tunnel {
        ldp-p2mp;
    }

    selective {
        group 225.0.1.0/24 {

```

```

        source 0.0.0.0/0 {
            ldp-p2mp;

            threshold-rate 0;
        }
    }
}

vrf-target target:100:10;
protocols {
    ospf {
        export parent_vpn_routes;
        area 0.0.0.0 {
            interface lo0.105 {
                passive;
            }
            interface ge-1/2/15.0;
        }
    }
    pim {
        rp {
            static {
                address 10.100.1.2;
            }
        }
        interface ge-1/2/15.0 {
            mode sparse;
        }
    }
    mvpn {
        mvpn-mode {
            rpt-spt;
        }
        sender-based-rpf;
    }
}
}
}

```

```

user@PE2# show routing-options
router-id 10.1.1.5;

```

```
route-distinguisher-id 10.1.1.5;
autonomous-system 65001;
```

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

Verification

IN THIS SECTION

- [Verifying Sender-Based RPF | 988](#)
- [Checking the BGP Routes | 989](#)
- [Checking the PIM Joins on the Downstream CE Receiver Devices | 997](#)
- [Checking the PIM Joins on the PE Devices | 998](#)
- [Checking the Multicast Routes | 1001](#)
- [Checking the MVPN C-Multicast Routes | 1003](#)
- [Checking the Source PE | 1004](#)

Confirm that the configuration is working properly.

Verifying Sender-Based RPF

Purpose

Make sure that sender-based RPF is enabled on Device PE2.

Action

```
user@PE2> show mvpn instance vpn-1

MVPN instance:
Legend for provider tunnel
S-   Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Family : INET
```

```

Instance : vpn-1
  MVPN Mode : RPT-SPT
  Sender-Based RPF: Enabled.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:LDP-P2MP:10.1.1.4, lsp-id 16777217
  Neighbor                               Inclusive Provider Tunnel
  10.1.1.2                               LDP-P2MP:10.1.1.2, lsp-id 16777219
  10.1.1.5                               LDP-P2MP:10.1.1.5, lsp-id 16777210
  C-mcast IPv4 (S:G)                    Provider Tunnel                St
  0.0.0.0/0:224.1.1.1/32                LDP-P2MP:10.1.1.2, lsp-id 16777219
  0.0.0.0/0:224.2.127.254/32           LDP-P2MP:10.1.1.3, lsp-id 16777210

MVPN instance:
Legend for provider tunnel
S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)           RM -- remote VPN route
Family : INET6

Instance : vpn-1
  MVPN Mode : RPT-SPT
  Sender-Based RPF: Enabled.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:LDP-P2MP:10.1.1.4, lsp-id 16777217

```

Checking the BGP Routes

Purpose

Make sure the expected BGP routes are being added to the routing tables on the PE devices.

Action

```

user@PE1> show route protocol bgp

inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

vpn-1.inet.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)

```

+ = Active Route, - = Last Active, * = Both

```

10.1.1.6/32      *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.7/32      *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.1.1.16/30     *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.20/30     *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.100.1.4/32    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.100.1.5/32    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299776(top)

```

vpn-1.inet.1: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```

10.1.1.4:32767:10.1.1.6/32
                 *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.1.1.16/30
                 *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.100.1.4/32
                 *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
                 > via ge-1/2/11.0, Push 299776, Push 299792(top)
10.1.1.5:32767:10.1.1.7/32
                 *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified

```



```

    > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.1.1.20/30
    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.100.1.5/32
    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776, Push 299776(top)

bgp.mvpn.0: 5 destinations, 8 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.4:32767:10.1.1.4/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.1.1.1/240
    *[BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 1d 04:17:24, MED 0, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.2.127.254/240
    *[BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 1d 04:17:23, MED 0, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
7:10.1.1.2:32767:1001:32:10.1.1.1:32:224.1.1.1/240
    *[BGP/170] 20:34:47, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 20:34:47, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792

```

```

vpn-1.mvpn.0: 7 destinations, 13 routes (7 active, 2 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.4:32767:10.1.1.4/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.1.1.1/240
    [BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 1d 04:17:24, MED 0, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
6:10.1.1.2:32767:1001:32:10.100.1.2:32:224.2.127.254/240
    [BGP/170] 1d 04:17:25, MED 0, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776
    [BGP/170] 1d 04:17:23, MED 0, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
7:10.1.1.2:32767:1001:32:10.10.1.1:32:224.1.1.1/240
    [BGP/170] 20:34:47, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299792
    [BGP/170] 20:34:47, localpref 100, from 10.1.1.5
    AS path: I, validation-state: unverified
    > via ge-1/2/11.0, Push 299776

```

```

user@PE2> show route protocol bgp

```

```

inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)

```

```

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

```

vpn-1.inet.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```

10.1.1.1/32      *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.7/32      *[BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299776(top)
10.1.1.0/30      *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.20/30     *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299776(top)
10.100.1.2/32    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299808(top)
10.100.1.5/32    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299776(top)

```

vpn-1.inet.1: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```

10.1.1.2:32767:10.1.1.1/32
                 *[BGP/170] 1d 04:23:24, MED 1, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.1.1.0/30
                 *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.100.1.2/32
                 *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
                 > via ge-1/2/12.0, Push 299776, Push 299808(top)
10.1.1.5:32767:10.1.1.7/32
                 *[BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.5

```

```

        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.1.1.20/30
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299776, Push 299776(top)
10.1.1.5:32767:10.100.1.5/32
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299776, Push 299776(top)

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

1:10.1.1.2:32767:10.1.1.2/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299808
1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299776
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
    *[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299808

vpn-1.mvpn.0: 7 destinations, 9 routes (7 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.2:32767:10.1.1.2/240
    *[BGP/170] 1d 04:23:24, localpref 100, from 10.1.1.2
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299808
1:10.1.1.5:32767:10.1.1.5/240
    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.5
        AS path: I, validation-state: unverified
        > via ge-1/2/12.0, Push 299776
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
    *[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
        AS path: I, validation-state: unverified

```

```
> via ge-1/2/12.0, Push 299808
```

```
user@PE3> show route protocol bgp
```

```
inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)
```

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
vpn-1.inet.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.1.1.1/32      *[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299808(top)
```

```
10.1.1.6/32      *[BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299792(top)
```

```
10.1.1.0/30      *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299808(top)
```

```
10.1.1.16/30     *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299792(top)
```

```
10.100.1.2/32    *[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299808(top)
```

```
10.100.1.4/32    *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
                 AS path: I, validation-state: unverified
```

```
> via ge-1/2/13.0, Push 299776, Push 299792(top)
```

```
vpn-1.inet.1: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
```

```
bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.1.1.2:32767:10.1.1.1/32
```

```
*[BGP/170] 1d 04:23:23, MED 1, localpref 100, from 10.1.1.2
                 AS path: I, validation-state: unverified
```

```

> via ge-1/2/13.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.1.1.0/30
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299808(top)
10.1.1.2:32767:10.100.1.2/32
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299808(top)
10.1.1.4:32767:10.1.1.6/32
*[BGP/170] 1d 04:23:20, MED 1, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.1.1.16/30
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299792(top)
10.1.1.4:32767:10.100.1.4/32
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299776, Push 299792(top)

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

1:10.1.1.2:32767:10.1.1.2/240
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299808
1:10.1.1.4:32767:10.1.1.4/240
*[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299792
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
*[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
AS path: I, validation-state: unverified
> via ge-1/2/13.0, Push 299808

vpn-1.mvpn.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.2:32767:10.1.1.2/240
*[BGP/170] 1d 04:23:23, localpref 100, from 10.1.1.2

```

```

                AS path: I, validation-state: unverified
                > via ge-1/2/13.0, Push 299808
1:10.1.1.4:32767:10.1.1.4/240
                *[BGP/170] 1d 04:23:20, localpref 100, from 10.1.1.4
                AS path: I, validation-state: unverified
                > via ge-1/2/13.0, Push 299792
5:10.1.1.2:32767:32:10.1.1.1:32:224.1.1.1/240
                *[BGP/170] 20:34:47, localpref 100, from 10.1.1.2
                AS path: I, validation-state: unverified
                > via ge-1/2/13.0, Push 299808

```

Checking the PIM Joins on the Downstream CE Receiver Devices

Purpose

Make sure that the expected join messages are being sent.

Action

```

user@CE2> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-1/2/14.0

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-1/2/14.0

Instance: PIM.master Family: INET6

```

```
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
-----
```

```
user@CE3> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-1/2/15.0

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-1/2/15.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
-----
```

Meaning

Both Device CE2 and Device CE3 send C-Join packets upstream to their neighboring PE routers, their unicast next-hop to reach the C-Source.

Checking the PIM Joins on the PE Devices

Purpose

Make sure that the expected join messages are being sent.

Action

```
user@PE1> show pim join instance vpn-1
Instance: PIM.vpn-1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```



```
Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 224.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream interface: ge-1/2/10.0

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
```

```
user@PE2> show pim join instance vpn-1
Instance: PIM.vpn-1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP

Group: 224.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream protocol: BGP
  Upstream interface: Through BGP

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
```

Upstream interface: Through BGP

```
user@PE3> show pim join instance vpn-1
Instance: PIM.vpn-1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP

Group: 224.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream protocol: BGP
  Upstream interface: Through BGP

Group: 224.2.127.254
  Source: *
  RP: 10.100.1.2
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP
```

Meaning

Both Device CE2 and Device CE3 send C-Join packets upstream to their neighboring PE routers, their unicast next-hop to reach the C-Source.

The C-Join state points to BGP as the upstream interface. Actually, there is no PIM neighbor relationship between the PEs. The downstream PE converts the C-PIM (C-S, C-G) state into a Type 7 source-tree join BGP route, and sends it to the upstream PE router toward the C-Source.

Checking the Multicast Routes

Purpose

Make sure that the C-Multicast flow is integrated in MVPN vpn-1 and sent by Device PE1 into the provider tunnel.

Action

```
user@PE1> show multicast route instance vpn-1
```

```
Instance: vpn-1 Family: INET
```

```
Group: 224.1.1.1/32
```

```
Source: *
```

```
Upstream interface: local
```

```
Downstream interface list:
```

```
ge-1/2/11.0
```

```
Group: 224.1.1.1
```

```
Source: 10.1.1.1/32
```

```
Upstream interface: ge-1/2/10.0
```

```
Downstream interface list:
```

```
ge-1/2/11.0
```

```
Group: 224.2.127.254/32
```

```
Source: *
```

```
Upstream interface: local
```

```
Downstream interface list:
```

```
ge-1/2/11.0
```

```
user@PE2> show multicast route instance vpn-1
```

```
Instance: vpn-1 Family: INET
```

```
Group: 224.1.1.1/32
```

```
Source: *
```

```
Upstream rpf interface list:
```

```
vt-1/2/10.4 (P)
```

```
Sender Id: Label 299840
```

```
Downstream interface list:
```

```
ge-1/2/14.0
```

```

Group: 224.1.1.1
  Source: 10.1.1.1/32
  Upstream rpf interface list:
    vt-1/2/10.4 (P)
      Sender Id: Label 299840

```

```

Group: 224.2.127.254/32
  Source: *
  Upstream rpf interface list:
    vt-1/2/10.4 (P)
      Sender Id: Label 299840
  Downstream interface list:
    ge-1/2/14.0

```

```

user@PE3> show multicast route instance vpn-1

```

```

Instance: vpn-1 Family: INET

```

```

Group: 224.1.1.1/32
  Source: *
  Upstream interface: vt-1/2/10.5
  Downstream interface list:
    ge-1/2/15.0

```

```

Group: 224.1.1.1
  Source: 10.1.1.1/32
  Upstream interface: vt-1/2/10.5

```

```

Group: 224.2.127.254/32
  Source: *
  Upstream interface: vt-1/2/10.5
  Downstream interface list:
    ge-1/2/15.0

```

Meaning

The output shows that, unlike the other PE devices, Device PE2 is using sender-based RPF. The output on Device PE2 includes the upstream RPF sender. The Sender Id field is only shown when sender-based RPF is enabled.

Checking the MVPN C-Multicast Routes

Purpose

Check the MVPN C-multicast route information,

Action

```
user@PE1> show mvpn c-multicast instance-name vpn-1
```

MVPN instance:

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET

Instance : vpn-1

MVPN Mode : RPT-SPT

C-mcast IPv4 (S:G)	Provider Tunnel	St	
0.0.0.0/0:224.1.1.1/32	I-P-tnl:LDP-P2MP:1.1.1.3, lsp-id 16777217		RM
10.1.1.1/32:224.1.1.1/32	I-P-tnl:LDP-P2MP:1.1.1.3, lsp-id 16777217		RM
0.0.0.0/0:224.2.127.254/32	I-P-tnl:LDP-P2MP:1.1.1.3, lsp-id 16777217		RM

...

```
user@PE2> show mvpn c-multicast instance-name vpn-1
```

MVPN instance:

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET

Instance : vpn-1

MVPN Mode : RPT-SPT

C-mcast IPv4 (S:G)	Provider Tunnel	St	

```

0.0.0.0/0:224.1.1.1/32      I-P-tnl:LDP-P2MP:1.1.1.2, lsp-id 16777217
10.1.1.1/32:224.1.1.1/32   I-P-tnl:LDP-P2MP:1.1.1.2, lsp-id 16777217
0.0.0.0/0:224.2.127.254/32 I-P-tnl:LDP-P2MP:1.1.1.2, lsp-id 16777217
...

```

```
user@PE3> show mvpn c-multicast instance-name vpn-1
```

MVPN instance:

Legend for provider tunnel

S- Selective provider tunnel

Legend for c-multicast routes properties (Pr)

DS -- derived from (*, c-g) RM -- remote VPN route

Family : INET

Instance : vpn-1

MVPN Mode : RPT-SPT

C-mcast IPv4 (S:G)	Provider Tunnel	St
0.0.0.0/0:224.1.1.1/32	I-P-tnl:LDP-P2MP:1.1.1.2, lsp-id 16777217	
10.1.1.1/32:224.1.1.1/32	I-P-tnl:LDP-P2MP:1.1.1.2, lsp-id 16777217	
0.0.0.0/0:224.2.127.254/32	I-P-tnl:LDP-P2MP:1.1.1.2, lsp-id 16777217	
...		

...

Meaning

The output shows the provider tunnel and label information.

Checking the Source PE

Purpose

Check the details of the source PE,

Action

```
user@PE1> show mvpn c-multicast source-pe
```

Instance : vpn-1

```

MVPN Mode : RPT-SPT
Family : INET
C-Multicast route address :0.0.0.0/0:224.1.1.1/32
  MVPN Source-PE1:
    extended-community: no-advertise target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: lo0.102 Index: -1610691384
  PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: lo0.102 Index: -1610691384
C-Multicast route address :10.1.1.1/32:224.1.1.1/32
  MVPN Source-PE1:
    extended-community: no-advertise target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: ge-1/2/10.0 Index: -1610691384
  PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: ge-1/2/10.0 Index: -1610691384
C-Multicast route address :0.0.0.0/0:224.2.127.254/32
  MVPN Source-PE1:
    extended-community: no-advertise target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: lo0.102 Index: -1610691384
  PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: lo0.102 Index: -1610691384

```

```

user@PE2> show mvpn c-multicast source-pe
Instance : vpn-1
MVPN Mode : RPT-SPT
Family : INET
C-Multicast route address :0.0.0.0/0:224.1.1.1/32

```

```

MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
C-Multicast route address :10.1.1.1/32:224.1.1.1/32
MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
C-Multicast route address :0.0.0.0/0:224.2.127.254/32
MVPN Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)
PIM Source-PE1:
    extended-community: target:10.1.1.2:72
    Route Distinguisher: 10.1.1.2:32767
    Autonomous system number: 65001
    Interface: (Null)

```

```
user@PE3> show mvpn c-multicast source-pe
```

```

Instance : vpn-1
MVPN Mode : RPT-SPT
Family : INET
C-Multicast route address :0.0.0.0/0:224.1.1.1/32
MVPN Source-PE1:

```



```

        extended-community: target:10.1.1.2:72
        Route Distinguisher: 10.1.1.2:32767
        Autonomous system number: 65001
        Interface: (Null)
    PIM Source-PE1:
        extended-community: target:10.1.1.2:72
        Route Distinguisher: 10.1.1.2:32767
        Autonomous system number: 65001
        Interface: (Null)
    C-Multicast route address :10.1.1.1/32:224.1.1.1/32
    MVPN Source-PE1:
        extended-community: target:10.1.1.2:72
        Route Distinguisher: 10.1.1.2:32767
        Autonomous system number: 65001
        Interface: (Null)
    PIM Source-PE1:
        extended-community: target:10.1.1.2:72
        Route Distinguisher: 10.1.1.2:32767
        Autonomous system number: 65001
        Interface: (Null)
    C-Multicast route address :0.0.0.0/0:224.2.127.254/32
    MVPN Source-PE1:
        extended-community: target:10.1.1.2:72
        Route Distinguisher: 10.1.1.2:32767
        Autonomous system number: 65001
        Interface: (Null)
    PIM Source-PE1:
        extended-community: target:10.1.1.2:72
        Route Distinguisher: 10.1.1.2:32767
        Autonomous system number: 65001
        Interface: (Null)
    ...

```

Meaning

The output shows the provider tunnel and label information.

RELATED DOCUMENTATION

| [unicast-umh-election](#)

Configuring MBGP MVPN Wildcards

IN THIS SECTION

- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN | 1008](#)
- [Configuring a Selective Provider Tunnel Using Wildcards | 1014](#)
- [Example: Configuring Selective Provider Tunnels Using Wildcards | 1015](#)

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

IN THIS SECTION

- [About S-PMSI | 1008](#)
- [Scenarios for Using Wildcard S-PMSI | 1010](#)
- [Types of Wildcard S-PMSI | 1011](#)
- [Differences Between Wildcard S-PMSI and \(S,G\) S-PMSI | 1011](#)
- [Wildcard \(*,*\) S-PMSI and PIM Dense Mode | 1012](#)
- [Wildcard \(*,*\) S-PMSI and PIM-BSR | 1012](#)
- [Wildcard Source and the 0.0.0.0/0 Source Prefix | 1013](#)

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

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

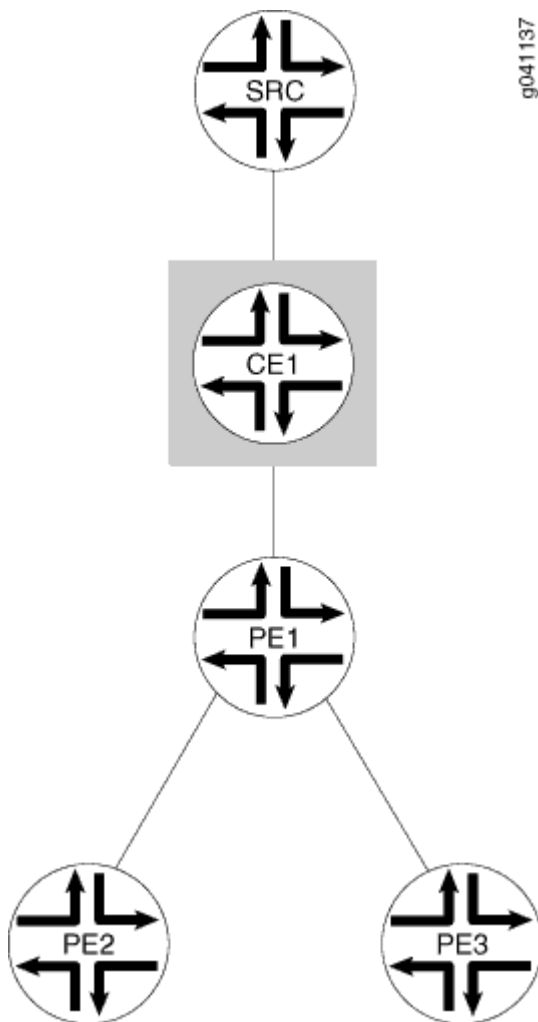
(S-PMSI) autodiscovery route advertises binding of a given MVPN customer multicast flow to a particular provider tunnel. The S-PMSI autodiscovery route advertised by the ingress PE router contains /32 IPv4 or /128 IPv6 addresses for the customer source and the customer group derived from the source-tree customer multicast route.

[Figure 119 on page 1010](#) 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 119: 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 {
                sptn13;
              }
            }
          }
        }
        wildcard-source {
          rsvp-te {
            label-switched-path-template {
              sptn12;
            }
            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 sptn13. A join message for (*, 203.0.113.0/24) is bound to sptn12.

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 <https://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 vjna 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 vjna 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 vjna 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 vjna 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
```

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 {
                sptn1;
              }
            }
          }
        }
      }
    }
    group 203.0.113.0/24 {
      wildcard-source {
        rsvp-te {
          label-switched-path-template {
            sptn2;
          }
        }
      }
    }
  }
}
```

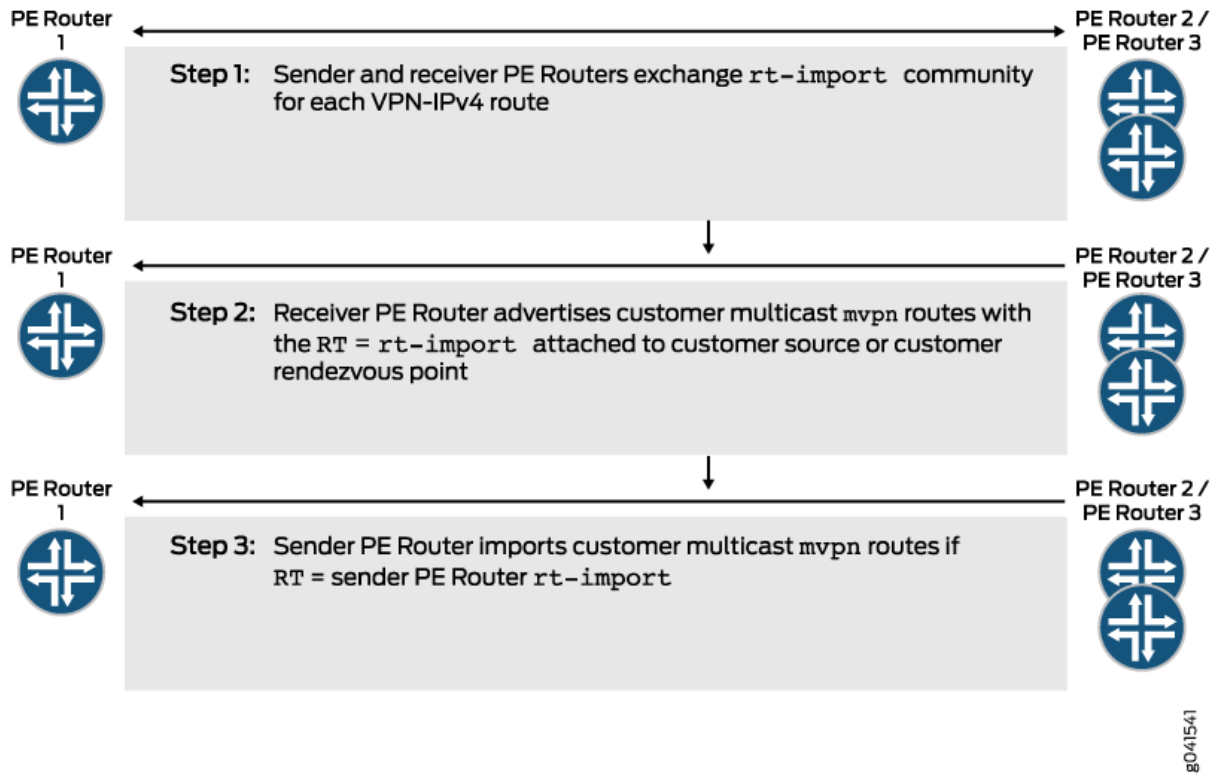

Distributing C-Multicast Routes Overview

IN THIS SECTION

- [Constructing C-Multicast Routes | 1019](#)
- [Eliminating PE-PE Distribution of \(C-*, C-G\) State Using Source Active Autodiscovery Routes | 1021](#)
- [Receiving C-Multicast Routes | 1022](#)

While non-C-multicast multicast virtual private network (MVPN) routes (Type 1 – Type 5) are generally used by all provider edge (PE) routers in the network, C-multicast MVPN routes (Type 6 and Type 7) are only useful to the PE router connected to the active C-S or candidate rendezvous point (RP). Therefore, C-multicast routes need to be installed only in the VPN routing and forwarding (VRF) table on the active sender PE router for a given C-G. To accomplish this, Internet draft [draft-ietf-l3vpn-2547bis-mcast-10.txt](#) specifies to attach a special and dynamic route target to C-multicast MVPN routes ([Figure 120 on page 1018](#)).

Figure 120: Attaching a Special and Dynamic Route Target to C-Multicast MVPN Routes



The route target attached to C-multicast routes is also referred to as the C-multicast import route target and should not to be confused with route target import (Table 32 on page 1018). Note that C-multicast MVPN routes differ from other MVPN routes in one essential way: they carry a dynamic route target whose value depends on the identity of the active sender PE router at a given time and can change if the active PE router changes.

Table 32: Distinction Between Route Target Import Attached to VPN-IPv4 Routes and Route Target Attached to C-Multicast MVPN Routes

Route Target Import Attached to VPN-IPv4 Routes	Route Target Attached to C-Multicast MVPN Routes
Value generated by the originating PE router. Must be unique per VRF table.	Value depends on the identity of the active PE router.

Table 32: Distinction Between Route Target Import Attached to VPN-IPv4 Routes and Route Target Attached to C-Multicast MVPN Routes (Continued)

Route Target Import Attached to VPN-IPv4 Routes	Route Target Attached to C-Multicast MVPN Routes
<p>Static. Created upon configuration to help identify to which PE router and to which VPN the VPN unicast routes belong.</p>	<p>Dynamic because if the active sender PE router changes, then the route target attached to the C-multicast routes must change to target the new sender PE router. For example, a new VPN source attached to a different PE router becomes active and preferred.</p>

A PE router that receives a local C-join determines the identity of the active sender PE router by performing a unicast route lookup for the C-S or candidate rendezvous point (router) [candidate RP] in the unicast VRF table. If there is more than one route, the receiver PE router chooses a single forwarder PE router. The procedures used for choosing a single forwarder are outlined in Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt and are not covered in this topic.

After the active sender (upstream) PE router is selected, the receiver PE router constructs the C-multicast MVPN route corresponding to the local C-join.

After the C-multicast route is constructed, the receiver PE router needs to attach the correct route target to this route targeting the active sender PE router. As mentioned, each PE router creates a unique VRF route target import community and attaches it to the VPN-IPv4 routes. When the receiver PE router does a route lookup for C-S or candidate RP, it can extract the value of the route target import associated with this route and set the value of the C-import route target to the value of the route target import.

On the active sender PE router, C-multicast routes are imported only if they carry the route target whose value is the same as the route target import that the sender PE router generated.

Constructing C-Multicast Routes

A PE router originates a C-multicast MVPN route in response to receiving a C-join through its PE-CE interface. See [Figure 121 on page 1020](#) for the formats in the C-multicast route encoded in MCAST-VPN NLRI. [Table 33 on page 1020](#) describes each field.

Figure 121: C-Multicast Route Type MCAST-VPN NLRI Format

Route Distinguisher	8 octets
Source AS	4 octets
Multicast Source Length	1 octet
Multicast Source	Variable
Multicast Group Length	1 octet
Multicast Group	Variable

8041542

Table 33: C-Multicast Route Type MCAST-VPN NLRI Format Descriptions

Field	Description
Route Distinguisher	Set to the route distinguisher of the C-S or candidate RP (the route distinguisher associated with the upstream PE router).
Source AS	Set to the value found in the src-as community of the C-S or candidate RP.
Multicast Source Length	Set to 32 for IPv4 and to 128 for IPv6 C-S or candidate RP IP addresses.
Multicast Source	Set to the IP address of the C-S or candidate RP.
Multicast Group Length	Set to 32 for IPv4 and to 128 for IPv6 C-G addresses.
Multicast Group	Set to the C-G of the received C-join.

This same structure is used for encoding both Type 6 and Type 7 routes with two differences:

- The first difference is the value used for the multicast source field. For Type 6 routes, this field is set to the IP address of the candidate RP configured. For Type 7 routes, this field is set to the IP address of the C-S contained in the (C-S, C-G) message.

- The second difference is the value used for the route distinguisher. For Type 6 routes, this field is set to the route distinguisher that is attached to the IP address of the candidate RP. For Type 7 routes, this field is set to the route distinguisher that is attached to the IP address of the C-S.

Eliminating PE-PE Distribution of (C-*, C-G) State Using Source Active Autodiscovery Routes

PE routers must maintain additional state when the C-multicast routing protocol is Protocol Independent Multicast-Sparse Mode (PIM-SM) in any-source multicast (ASM). This is a requirement because with ASM, the receivers first join the shared tree rooted at the candidate RP (called a candidate RP tree or candidate RPT). However, as the VPN multicast sources become active, receivers learn the identity of the sources and join the tree rooted at the source (called a customer shortest-path tree or C-SPT). The receivers then send a prune message to the candidate RP to stop the traffic coming through the shared tree for the group that they joined to the C-SPT. The switch from candidate RPT to C-SPT is a complicated process requiring additional state.

Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt specifies optional procedures that completely eliminate the need for joining the candidate RPT. These procedures require PE routers to keep track of all active VPN sources using one of two options. The first option is to colocate the candidate RP on one of the PE routers. The second option is to use the Multicast Source Discovery Protocol (MSDP) between one of the PE routers and the customer candidate RP.

In this approach, a PE router that receives a local (C-*, C-G) join creates a Type 6 route, but does not advertise the route to the remote PE routers until it receives information about an active source. The PE router acting as the candidate RP (or that learns about active sources via MSDP) is responsible for originating a Type 5 route. A Type 5 route carries information about the active source and the group addresses. The information contained in a Type 5 route is enough for receiver PE routers to join the C-SPT by originating a Type 7 route toward the sender PE router, completely skipping the advertisement of the Type 6 route that is created when a C-join is received. [Figure 122 on page 1022](#) shows the format of a source active (SA) autodiscovery route. [Table 34 on page 1022](#) describes each format.

Figure 122: Source Active Autodiscovery Route Type MCAST-VPN NLRI Format

Route Distinguisher	8 octets
Multicast Source Length	1 octet
Multicast Source	Variable
Multicast Group Length	1 octet
Multicast Group	Variable

8041543

Table 34: Source Active Autodiscovery Route Type MCAST-VPN NLRI Format Descriptions

Field	Description
Route Distinguisher	Set to the route distinguisher configured on the router originating the SA autodiscovery route.
Multicast Source Length	Set to 32 for IPv4 and to 128 for IPv6 C-S IP addresses.
Multicast Source	Set to the IP address of the C-S that is actively transmitting data to C-G.
Multicast Group Length	Set to 32 for IPv4 and to 128 for IPv6 C-G addresses.
Multicast Group	Set to the IP address of the C-G to which C-S is transmitting data.

Receiving C-Multicast Routes

The sender PE router imports C-multicast routes into the VRF table based on the route target of the route. If the route target attached to the C-multicast MVPN route matches the route target import community originated by this router, the C-multicast MVPN route is imported into the VRF table. If not, it is discarded.

Once the C-multicast MVPN routes are imported, they are translated back to C-joins and passed on to the VRF C-PIM protocol for further processing per normal PIM procedures.

RELATED DOCUMENTATION

[Enabling Next-Generation MVPN Services | 734](#)

[Exchanging C-Multicast Routes | 1023](#)

[Understanding Next-Generation MVPN Network Topology | 716](#)

Exchanging C-Multicast Routes

IN THIS SECTION

- [Advertising C-Multicast Routes Using BGP | 1023](#)
- [Receiving C-Multicast Routes | 1029](#)

This section describes PE-PE distribution of Type 7 routes discussed in "[Signaling Provider Tunnels and Data Plane Setup](#)" on page 1037.

In source-tree-only mode, a receiver provider edge (PE) router generates and installs a Type 6 route in its `<routing-instance-name>.mvpn.0` table in response to receiving a (C-*, C-G) message from a local receiver, but does not advertise this route to other PE routers via BGP. The receiver PE router waits for a Type 5 route corresponding to the C-join.

Type 5 routes carry information about active sources and can be advertised by any PE router. In Junos OS, a PE router originates a Type 5 route if one of the following conditions occurs:

- PE router starts receiving multicast data directly from a VPN multicast source.
- PE router is the candidate rendezvous point (router) (candidate RP) and starts receiving C-PIM register messages.
- PE router has a Multicast Source Discovery Protocol (MSDP) session with the candidate RP and starts receiving MSDP Source Active routes.

Once both Type 6 and Type 5 routes are installed in the `<routing-instance-name>.mvpn.0` table, the receiver PE router is ready to originate a Type 7 route

Advertising C-Multicast Routes Using BGP

If the C-join received over a VPN interface is a source tree join (C-S, C-G), then the receiver PE router simply originates a Type 7 route (Step 7 in the following procedure). If the C-join is a shared tree join (C-

*, C-G), then the receiver PE router needs to go through a few steps (Steps 1-7) before originating a Type 7 route.

Note that Router PE1 is the candidate RP that is conveniently located in the same router as the sender PE router. If the sender PE router and the PE router acting as (or MSDP peering with) the candidate RP are different, then the VPN multicast register messages first need to be delivered to the PE router acting as the candidate RP that is responsible for originating the Type 5 route. Routers referenced in this topic are shown in "[Understanding Next-Generation MVPN Network Topology](#)" on page 716.

1. A PE router that receives a (C-*, C-G) join message processes the message using normal C-PIM procedures and updates its C-PIM database accordingly.

Enter the `show pim join extensive instance vpna 224.1.1.1` command on Router PE3 to verify that Router PE3 creates the C-PIM database after receiving the (*, 224.1.1.1) C-join message from Router CE3:

```
user@PE3> show pim join extensive instance vpna
224.1.1.1
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.12.53.1
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP
  Upstream neighbor: Through MVPN
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: so-0/2/0.0
      10.12.87.1 State: Join Flags: SRW Timeout: Infinity
```

2. The (C-*, C-G) entry in the C-PIM database triggers the generation of a Type 6 route that is then installed in the `<routing-instance-name>.mvpn.0` table by C-PIM. The Type 6 route uses the candidate RP IP address as the source.

Enter the `show route table vpna.mvpn.0 detail | find 6:10.1.1.1` command on Router PE3 to verify that Router PE3 installs the following Type 6 route in the `vpna.mvpn.0` table:

```
user@PE3> show route table vpna.mvpn.0 detail
| find 6:10.1.1.1
6:10.1.1.1:1:65000:32:10.12.53.1:32:224.1.1.1/240 (1 entry, 1 announced)
      *PIM      Preference: 105
```

```

Next hop type: Multicast (IPv4), Next hop index: 262144
Next-hop reference count: 11
State: <Active Int>
Age: 1d 1:32:58
Task: PIM.vpna
Announcement bits (2): 0-PIM.vpna 1-mvpn global task
AS path: I
Communities: no-advertise target:10.1.1.1:64

```

3. The route distinguisher and route target attached to the Type 6 route are learned from a route lookup in the <routing-instance-name>.inet.0 table for the IP address of the candidate RP.

Enter the show route table vpna.inet.0 10.12.53.1 detail command on Router PE3 to verify that Router PE3 has the following entry for C-RP 10.12.53.1 in the vpna.inet.0 table:

```

user@PE3> show route table vpna.inet.0 10.12.53.1
detail
vpna.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
10.12.53.1/32 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
              Route Distinguisher: 10.1.1.1:1
              Next hop type: Indirect
              Next-hop reference count: 6
              Source: 10.1.1.1
              Next hop type: Router, Next hop index: 588
              Next hop: via so-0/0/3.0, selected
              Label operation: Push 16, Push 299808(top)
              Protocol next hop: 10.1.1.1
              Push 16
              Indirect next hop: 8da91f8 262143
              State: <Secondary Active Int Ext>
              Local AS: 65000 Peer AS: 65000
              Age: 4:49:25 Metric2: 1
              Task: BGP_65000.10.1.1.1+179
              Announcement bits (1): 0-KRT
              AS path: I
              Communities: target:10:1 src-as:65000:0 rt-import:10.1.1.1:64
              Import Accepted
              VPN Label: 16
              Localpref: 100

```

```
Router ID: 10.1.1.1
Primary Routing Table bgp.13vpn.0
```

4. After the VPN source starts transmitting data, the first PE router that becomes aware of the active source (either by receiving register messages or the MSDP source-active routes) installs a Type 5 route in its VRF `mvpn` table.

Enter the `show route table vpna.mvpn.0 detail | find 5:10.1.1.1` command on Router PE1 to verify that Router PE1 has installed the following entry in the `vpna.mvpn.0` table and starts receiving C-PIM register messages from Router CE1:

```
user@PE1> show route table vpna.mvpn.0 detail
| find 5:10.1.1.1
5:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
    *PIM      Preference: 105
             Next hop type: Multicast (IPv4)
             Next-hop reference count: 30
             State: <Active Int>
             Age: 1d 1:36:33
             Task: PIM.vpna
             Announcement bits (3): 0-PIM.vpna 1-mvpn global task 2-BGP RT Background
             AS path: I
```

5. Type 5 routes that are installed in the `<routing-instance-name>.mvpn.0` table are picked up by BGP and advertised to remote PE routers.

Enter the `show route advertising-protocol bgp 10.1.1.3 detail table vpna.mvpn.0 | find 5:` command on Router PE1 to verify that Router PE1 advertises the following Type 5 route to remote PE routers:

```
user@PE1> show route advertising-protocol bgp
10.1.1.3 detail table vpna.mvpn.0 | find 5:
* 5:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
    BGP group int type Internal
    Route Distinguisher: 10.1.1.1:1
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:10:1
```

6. The receiver PE router that has both a Type 5 and Type 6 route for (C-*, C-G) is now ready to originate a Type 7 route.

Enter the `show route table vpna.mvpn.0 detail` command on Router PE3 to verify that Router PE3 has the following Type 5, 6, and 7 routes in the `vpna.mvpn.0` table.

The Type 6 route is installed by C-PIM in Step 2. The Type 5 route is learned via BGP in Step 5. The Type 7 route is originated by the MVPN module in response to having both Type 5 and Type 6 routes for the same (C-*, C-G). The route target of the Type 7 route is the same as the route target of the Type 6 route because both routes (IP address of the candidate RP [10.12.53.1] and the address of the VPN multicast source [192.168.1.2]) are reachable via the same router [PE1]). Therefore, 10.12.53.1 and 192.168.1.2 carry the same route target import (10.1.1.1:64) community

```

user@PE3> show route table vpna.mvpn.0 detail
5:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Next hop type: Indirect
              Next-hop reference count: 4
              Source: 10.1.1.1
              Protocol next hop: 10.1.1.1
              Indirect next hop: 2 no-forward
              State: <Secondary Active Int Ext>
              Local AS: 65000 Peer AS: 65000
              Age: 1d 1:43:13 Metric2: 1
              Task: BGP_65000.10.1.1.1+55384
              Announcement bits (2): 0-PIM.vpna 1-mvpn global task
              AS path: I
              Communities: target:10:1
              Import Accepted
              Localpref: 100
              Router ID: 10.1.1.1
              Primary Routing Table bgp.mvpn.0

6:10.1.1.1:1:65000:32:10.12.53.1:32:224.1.1.1/240 (1 entry, 1 announced)
    *PIM      Preference: 105
              Next hop type: Multicast (IPv4), Next hop index: 262144
              Next-hop reference count: 11
              State: <Active Int>
              Age: 1d 1:44:09
              Task: PIM.vpna
              Announcement bits (2): 0-PIM.vpna 1-mvpn global task
              AS path: I
              Communities: no-advertise target:10.1.1.1:64

```

```

7:10.1.1.1:1:65000:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
  *MVPN      Preference: 70
              Next hop type: Multicast (IPv4), Next hop index: 262144
              Next-hop reference count: 11
              State: <Active Int Ext>
              Age: 1d 1:44:09 Metric2: 1
              Task: mvpn global task
              Announcement bits (3): 0-PIM.vpna 1-mvpn global task 2-BGP RT Background
              AS path: I
              Communities: target:10.1.1.1:64

```

7. The Type 7 route installed in the VRF MVPN table is picked up by BGP and advertised to remote PE routers.

Enter the `show route advertising-protocol bgp 10.1.1.1 detail table vpna.mvpn.0 | find 7:10.1.1.1` command on Router PE3 to verify that Router PE3 advertises the following Type 7 route:

```

user@PE3> show route advertising-protocol bgp
10.1.1.1 detail table vpna.mvpn.0 | find 7:10.1.1.1
* 7:10.1.1.1:1:65000:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
  BGP group int type Internal
  Route Distinguisher: 10.1.1.3:1
  Nexthop: Self
  Flags: Nexthop Change
  Localpref: 100
  AS path: [65000] I
  Communities: target:10.1.1.1:64

```

8. If the C-join is a source tree join, then the Type 7 route is originated immediately (without waiting for a Type 5 route).

Enter the `show route table vpna.mvpn.0 detail | find 7:10.1.1.1` command on Router PE2 to verify that Router PE2 originates the following Type 7 route in response to receiving a (192.168.1.2, 232.1.1.1) C-join:

```

user@PE2> show route table vpna.mvpn.0 detail
| find 7:10.1.1.1
7:10.1.1.1:1:65000:32:192.168.1.2:32:232.1.1.1/240 (1 entry, 1 announced)
  *PIM      Preference: 105
              Next hop type: Multicast (IPv4), Next hop index: 262146

```

```

Next-hop reference count: 4
State: <Active Int>
Age: 2d 18:59:56
Task: PIM.vpna
Announcement bits (3): 0-PIM.vpna 1-mvpn global task 2-BGP RT Background
AS path: I
Communities: target:10.1.1.1:64

```

Receiving C-Multicast Routes

A sender PE router imports a Type 7 route if the route is carrying a route target that matches the locally originated route target import community. All Type 7 routes must pass the `__vrf-mvpn-import-cmcast-<routing-instance-name>-internal__` policy in order to be installed in the `<routing-instance-name>.mvpn.0` table.

When a sender PE router receives a Type 7 route via BGP, this route is installed in the `<routing-instance-name>.mvpn.0` table. The BGP route is then translated back into a normal C-join inside the VRF table, and the C-join is installed in the local C-PIM database of the receiver PE router. A new C-join added to the C-PIM database triggers C-PIM to originate a Type 6 or Type 7 route. The C-PIM on the sender PE router creates its own version of the same Type 7 route received via BGP.

Use the `show route table vpna.mvpn.0 detail | find 7:10.1.1.1` command to verify that Router PE1 contains the following entries for a Type 7 route in the `vpna.mvpn.0` table corresponding to a (192.168.1.2, 224.1.1.1) join message. There are two entries; one entry is installed by PIM and the other entry is installed by BGP. This example also shows the Type 7 route corresponding to the (192.168.1.2, 232.1.1.1) join.

```

user@PE1> show route table vpna.mvpn.0 detail
| find 7:10.1.1.1
7:10.1.1.1:1:65000:32:192.168.1.2:32:224.1.1.1/240 (2 entries, 2 announced)
 *PIM      Preference: 105
           Next hop type: Multicast (IPv4)
           Next-hop reference count: 30
           State: <Active Int>
           Age: 1d 2:19:04
           Task: PIM.vpna
           Announcement bits (2): 0-PIM.vpna 1-mvpn global task
           AS path: I
           Communities: no-advertise target:10.1.1.1:64
BGP       Preference: 170/-101
           Next hop type: Indirect
           Next-hop reference count: 4
           Source: 10.1.1.3
           Protocol next hop: 10.1.1.3

```

```

Indirect next hop: 2 no-forward
State: <Secondary Int Ext>
Inactive reason: Route Preference
Local AS: 65000 Peer AS: 65000
Age: 53:27      Metric2: 1
Task: BGP_65000.10.1.1.3+179
Announcement bits (2): 0-PIM.vpna 1-mvpn global task
AS path: I
Communities: target:10.1.1.1:64
Import Accepted
Localpref: 100
Router ID: 10.1.1.3
Primary Routing Table bgp.mvpn.0
7:10.1.1.1:1:65000:32:192.168.1.2:32:232.1.1.1/240 (2 entries, 2 announced)
*BGP      Preference: 105
          Next hop type: Multicast (IPv4)
          Next-hop reference count: 30
          State: <Active Int>
          Age: 2d 19:21:17
          Task: PIM.vpna
          Announcement bits (2): 0-PIM.vpna 1-mvpn global task
          AS path: I
          Communities: no-advertise target:10.1.1.1:64
*BGP      Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 4
          Source: 10.1.1.2
          Protocol next hop: 10.1.1.2
          Indirect next hop: 2 no-forward
          State: <Secondary Int Ext>
          Inactive reason: Route Preference
          Local AS: 65000 Peer AS: 65000
          Age: 53:27      Metric2: 1
          Task: BGP_65000.10.1.1.2+49165
          Announcement bits (2): 0-PIM.vpna 1-mvpn global task
          AS path: I
          Communities: target:10.1.1.1:64
          Import Accepted
          Localpref: 100
          Router ID: 10.1.1.2
          Primary Routing Table bgp.mvpn.0

```


Remote C-joins (Type 7 routes learned via BGP translated back to normal C-joins) are installed in the VRF C-PIM database on the sender PE router and are processed based on regular C-PIM procedures. This process completes the end-to-end C-multicast routing exchange.

Use the `show pim join extensive instance vpna` command to verify that Router PE1 has installed the following entries in the C-PIM database:

```
user@PE1> show pim join extensive instance vpna
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: 192.168.1.2
  Flags: sparse,spt
  Upstream interface: fe-0/2/0.0
  Upstream neighbor: 10.12.97.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 201
  Downstream neighbors:
  Interface: Pseudo-MVPN

Group: 232.1.1.1
  Source: 192.168.1.2
  Flags: sparse,spt
  Upstream interface: fe-0/2/0.0
  Upstream neighbor: 10.12.97.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout:
  Downstream neighbors:
  Interface: Pseudo-MVPN

Instance: PIM.vpna Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

RELATED DOCUMENTATION

[Signaling Provider Tunnels and Data Plane Setup | 1037](#)

[Distributing C-Multicast Routes Overview | 1017](#)

[Understanding MBGP Multicast VPN Extranets | 858](#)

Generating Source AS and Route Target Import Communities Overview

Both route target import (`rt-import`) and source autonomous system (`src-as`) communities contain two fields (following their respective keywords). In Junos OS, a provider edge (PE) router constructs the route target import community using its router ID in the first field and a per-VRF unique number in the second field. The router ID is normally set to the primary loopback IP address of the PE router. The unique number used in the second field is an internal number derived from the routing-instance table index. The combination of the two numbers creates a route target import community that is unique to the originating PE router and unique to the VPN routing and forwarding (VRF) instance from which it is created.

For example, Router PE1 creates the following route target import community: `rt-import:10.1.1.1:64`.

Since the route target import community is constructed using the primary loopback address and the routing-instance table index of the PE router, any event that causes either number to change triggers a change in the value of the route target import community. This in turn requires VPN-IPv4 routes to be re-advertised with the new route target import community. Under normal circumstances, the primary loopback address and the routing-instance table index numbers do not change. If they do change, Junos OS updates all related internal policies and re-advertises VPN-IPv4 routes with the new `rt-import` and `src-as` values per those policies.

To ensure that the route target import community generated by a PE router is unique across VRF tables, the Junos OS Policy module restricts the use of primary loopback addresses to next-generation multicast virtual private network (MVPN) internal policies only. You are not permitted to configure a route target for any VRF table (MVPN or otherwise) using the primary loopback address. The commit fails with an error if the system finds a user-configured route target that contains the IP address used in constructing the route target import community.

The global administrator field of the `src-as` community is set to the local AS number of the PE router originating the community, and the local administrator field is set to `0`. This community is used for inter-AS operations but needs to be carried along with all VPN-IPv4 routes.

For example, Router PE1 creates an `src-as` community with a value of `src-as:65000:0`.

RELATED DOCUMENTATION

[Originating Type 1 Intra-AS Autodiscovery Routes Overview | 1033](#)

[Generating Next-Generation MVPN VRF Import and Export Policies Overview | 737](#)

[Enabling Next-Generation MVPN Services | 734](#)

Originating Type 1 Intra-AS Autodiscovery Routes Overview

IN THIS SECTION

- [Attaching Route Target Community to Type 1 Routes | 1033](#)
- [Attaching the PMSI Attribute to Type 1 Routes | 1034](#)
- [Sender-Only and Receiver-Only Sites | 1037](#)

Every provider edge (PE) router that is participating in the next-generation multicast virtual private network (MVPN) is required to originate a Type 1 intra-AS autodiscovery route. In Junos OS, the MVPN module is responsible for installing the intra-AS autodiscovery route in the local `<routing-instance-name>.mvpn.0` table. All PE routers advertise their local Type 1 routes to each other. Routers referenced in this topic are shown in "[Understanding Next-Generation MVPN Network Topology](#)" on page 716.

Use the `show route table vjna.mvpn.0` command to verify that Router PE1 has installed intra-AS AD routes in the `vjna.mvpn.0` table. The route is installed by the MVPN protocol (meaning it is the MVPN module that originated the route), and the mask for the entire route is `/240`.

```
user@PE1> show route table vjna.mvpn.0
vjna.mvpn.0: 6 destinations, 9 routes (6 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.1:1:10.1.1.1/240
      *[MVPN/70] 04:09:44, metric2 1
      Indirect
```

Attaching Route Target Community to Type 1 Routes

Intra-AS AD routes are picked up by the BGP protocol from the `<routing-instance-name>.mvpn.0` table and advertised to the remote PE routers via the MCAST-VPN address family. By default, intra-AS autodiscovery routes carry the same route target community that is attached to the unicast VPN-IPv4 routes. If the unicast and multicast network topologies are not congruent, then you can configure a different set of import route target and export route target communities for non-C-multicast MVPN routes (C-multicast MVPN routes always carry a dynamic import route target).

Multicast route targets are configured by including the `import-target` and `export-target` statements at the `[edit routing-instances routing-instance-name protocols mvpn route-target]` hierarchy level.

Junos OS creates two additional internal policies in response to configuring multicast route targets. These policies are applied to non-C-multicast MVPN routes during import and export decisions. Multicast VPN routing and forwarding (VRF) internal import and export policies follow a naming convention similar to unicast VRF import and export policies. The contents of these policies are also similar to policies applied to unicast VPN routes.

The following list identifies the default policy names and where they are applied:

Multicast VRF import policy: `__vrf-mvpn-import-target-<routing-instance-name>-internal__`

Multicast VRF export policy: `__vrf-mvpn-export-target-<routing-instance-name>-internal__`

Use the `show policy __vrf-mvpn-import-target-vpna-internal__` command on Router PE1 to verify that Router PE1 has created the following internal MVPN policies if `import-target` and `export-target` are configured to be `target:10:2`:

```
user@PE1> show policy __vrf-mvpn-import-target-vpna-internal__
Policy __vrf-mvpn-import-target-vpna-internal__:
  Term unnamed:
    from community __vrf-mvpn-community-import-vpna-internal__ [target:10:2 ]
    then accept
  Term unnamed:
    then reject
user@PE1> show policy __vrf-mvpn-export-target-vpna-internal__
Policy __vrf-mvpn-export-target-vpna-internal__:
  Term unnamed:
    then community + __vrf-mvpn-community-export-vpna-internal__ [target:10:2 ] accept
```

The values in this example are as follows:

- Multicast import RT community: `__vrf-mvpn-community-import-vpna-internal__`
- Multicast export RT community: `__vrf-mvpn-community-export-vpna-internal__` Value: `target:10:2`

Attaching the PMSI Attribute to Type 1 Routes

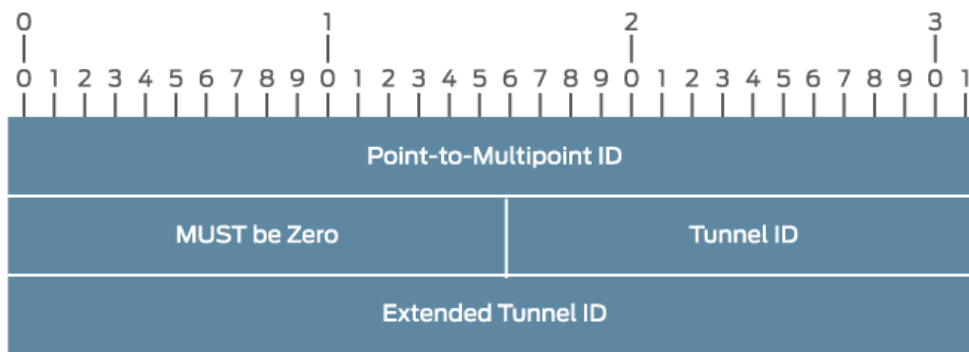
The provider multicast service interface (PMSI) attribute is originated and attached to Type 1 intra-AS autodiscovery routes by the sender PE routers when the `provider-tunnel` statement is included at the `[edit routing-instances <routing-instance-name>]` hierarchy level. Since provider tunnels are signaled by the sender PE routers, this statement is not necessary on the PE routers that are known to have VPN multicast receivers only.

If the provider tunnel configured is Protocol Independent Multicast-Sparse Mode (PIM-SM) any-source multicast (ASM), then the PMSI attribute carries the IP address of the sender-PE and provider tunnel

group address. The provider tunnel group address is assigned by the service provider (through configuration) from the provider's multicast address space and is not to be confused with the multicast addresses used by the VPN customer.

If the provider tunnel configured is the RSVP-Traffic Engineering (RSVP-TE) type, then the PMSI attribute carries the RSVP-TE point-to-multipoint session object. This point-to-multipoint session object is used as the identifier for the parent point-to-multipoint label-switched path (LSP) and contains the fields shown in [Figure 123 on page 1035](#).

Figure 123: RSVP-TE Point-to-Multipoint Session Object Format



8041547

In Junos OS, the P2MP ID and Extended Tunnel ID fields are set to the router ID of the sender PE router. The Tunnel ID is set to the port number used for the point-to-multipoint RSVP session that is unique for the length of the RSVP session.

Use the `show rsvp session p2mp detail` command to verify that Router PE1 signals the following RSVP sessions to Router PE2 and Router PE3 (using port number 6574). In this example, Router PE1 is signaling a point-to-multipoint LSP named `10.1.1.1:65535:mvpn:vpna` with two sub-LSPs. Both sub-LSPs `10.1.1.3:10.1.1.1:65535:mvpn:vpna` and `10.1.1.2:10.1.1.1:65535:mvpn:vpna` use the same RSVP port number (6574) as the parent point-to-multipoint LSP.

```

user@PE1> show rsvp session p2mp detail
Ingress RSVP: 2 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 2

10.1.1.3
  From: 10.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: 10.1.1.3:10.1.1.1:65535:mvpn:vpna, LSPpath: Primary
  P2MP LSPname: 10.1.1.1:65535:mvpn:vpna
  Suggested label received: -, Suggested label sent: -

```

Recovery label received: -, Recovery label sent: 299968
Resv style: 1 SE, Label in: -, Label out: 299968
Time left: -, Since: Wed May 27 07:36:22 2009
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 6574 protocol 0
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.12.100.6 (fe-0/2/3.0) 27 pkts
RESV rcvfrom: 10.12.100.6 (fe-0/2/3.0) 27 pkts
Explct route: 10.12.100.6 10.12.100.22
Record route: <self> 10.12.100.6 10.12.100.22

10.1.1.2

From: 10.1.1.1, LSPstate: Up, ActiveRoute: 0
LSPname: 10.1.1.2:10.1.1.1:65535:mvpn:vpna, LSPpath: Primary
P2MP LSPname: 10.1.1.1:65535:mvpn:vpna
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299968
Resv style: 1 SE, Label in: -, Label out: 299968
Time left: -, Since: Wed May 27 07:36:22 2009
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 6574 protocol 0
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.12.100.6 (fe-0/2/3.0) 27 pkts
RESV rcvfrom: 10.12.100.6 (fe-0/2/3.0) 27 pkts
Explct route: 10.12.100.6 10.12.100.9
Record route: <self> 10.12.100.6 10.12.100.9

Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Sender-Only and Receiver-Only Sites

In Junos OS, you can configure a PE router to be a sender-site only or a receiver-site only. These options are enabled by including the `sender-site` and `receiver-site` statements at the `[edit routing-instances routing-instance-name protocols mvpn]` hierarchy level.

- A sender-site only PE router does not join the provider tunnels advertised by remote PE routers
- A receiver-site only PE router does not send a PMSI attribute

The commit fails if you include the `receiver-site` and `provider-tunnel` statements in the same VPN.

RELATED DOCUMENTATION

[Generating Source AS and Route Target Import Communities Overview | 1032](#)

[Understanding MBGP Multicast VPN Extranets | 858](#)

[Signaling Provider Tunnels and Data Plane Setup | 1037](#)

[Generating Next-Generation MVPN VRF Import and Export Policies Overview | 737](#)

Signaling Provider Tunnels and Data Plane Setup

IN THIS SECTION

- [Provider Tunnels Signaled by PIM \(Inclusive\) | 1038](#)
- [Provider Tunnels Signaled by RSVP-TE \(Inclusive and Selective\) | 1043](#)

In a next-generation multicast virtual private network (MVPN), provider tunnel information is communicated to the receiver PE routers in an out-of-band manner. This information is advertised via BGP and is independent of the actual tunnel signaling process. Once the tunnel is signaled, the sender PE router binds the VPN routing and forwarding (VRF) table to the locally configured tunnel. The receiver PE routers bind the tunnel signaled to the VRF table where the Type 1 autodiscovery route with the matching provider multicast service interface (PMSI) attribute is installed. The same binding process is used for both Protocol Independent Multicast (PIM) and RSVP-Traffic Engineering (RSVP-TE) signaled provider tunnels.

Provider Tunnels Signaled by PIM (Inclusive)

A sender provider edge (PE) router configured to use an inclusive PIM-sparse mode (PIM-SM) any-source multicast (ASM) provider tunnel for a VPN creates a multicast tree (using the P-group address configured) in the service provider network. This tree is rooted at the sender PE router and has the receiver PE routers as the leaves. VPN multicast packets received from the local VPN source are encapsulated by the sender PE router with a multicast generic routing encapsulation (GRE) header containing the P-group address configured for the VPN. These packets are then forwarded on the service provider network as normal IP multicast packets per normal P-PIM procedures. At the leaf nodes, the GRE header is stripped and the packets are passed on to the local VRF C-PIM protocol for further processing.

In Junos OS, a logical interface called multicast tunnel (MT) is used for GRE encapsulation and de-encapsulation of VPN multicast packets. The multicast tunnel interface is created automatically if a Tunnel PIC is present.

- Encapsulation subinterfaces are created from an mt-x/y/z.[32768-49151] range.
- De-encapsulation subinterfaces are created from an mt-x/y/z.[49152-65535] range.

The multicast tunnel subinterfaces act as pseudo upstream or downstream interfaces between C-PIM and P-PIM.

In the following two examples, assume that the network uses PIM-SM (ASM) signaled GRE tunnels as the tunneling technology. Routers referenced in this topic are shown in "[Understanding Next-Generation MVPN Network Topology](#)" on page 716.

Use the `show interfaces mt-0/1/0 terse` command to verify that Router PE1 has created the following multicast tunnel subinterface. The logical interface number is 32768, indicating that this sub-unit is used for GRE encapsulation.

```
user@PE1> show interfaces mt-0/1/0 terse
```

Interface	Admin	Link	Proto	Local	Remote
mt-0/1/0	up	up			
mt-0/1/0.32768	up	up	inet		
			inet6		

Use the `show interfaces mt-0/1/0 terse` command to verify that Router PE2 has created the following multicast tunnel subinterface. The logical interface number is 49152, indicating that this sub-unit is used for GRE de-encapsulation.

```
user@PE2> show interfaces mt-0/1/0 terse
```

Interface	Admin	Link	Proto	Local	Remote


```

mt-0/1/0          up      up
mt-0/1/0.49152   up      up      inet
                                     inet6

```

P-PIM and C-PIM on the Sender PE Router

The sender PE router installs a local join entry in its P-PIM database for each VRF table configured to use PIM as the provider tunnel. The outgoing interface list (OIL) of this entry points to the core-facing interface. Since the P-PIM entry is installed as `Local`, the sender PE router sets the source address to its primary loopback IP address.

Use the `show pim join extensive` command to verify that Router PE1 has installed the following state in its P-PIM database.

```

user@PE1> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source
  Keepalive timeout: 339
  Downstream neighbors:
    Interface: fe-0/2/3.0
      10.12.100.6 State: Join Flags: S Timeout: 195

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

On the VRF side of the sender PE router, C-PIM installs a `Local Source` entry in its C-PIM database for the active local VPN source. The OIL of this entry points to `Pseudo-MVPN`, indicating that the downstream interface points to the receivers in the next-generation MVPN network. Routers referenced in this topic are shown in "[Understanding Next-Generation MVPN Network Topology](#)" on page 716.

Use the `show pim join extensive instance vpna 224.1.1.1` command to verify that Router PE1 has installed the following entry in its C-PIM database.

```

user@PE1> show pim join extensive instance vpna
224.1.1.1
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: 192.168.1.2
  Flags: sparse,spt
  Upstream interface: fe-0/2/0.0
  Upstream neighbor: 10.12.97.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 0
  Downstream neighbors:
    Interface: Pseudo-MVPN

```

The forwarding entry corresponding to the C-PIM Local Source (or Local RP) on the sender PE router points to the multicast tunnel encapsulation subinterface as the downstream interface. This indicates that the local multicast data packets are encapsulated as they are passed on to the P-PIM protocol.

Use the `show multicast route extensive instance vpna group 224.1.1.1` command to verify that Router PE1 has the following multicast forwarding entry for group 224.1.1.1. The upstream interface is the PE-CE interface and the downstream interface is the multicast tunnel encapsulation subinterface:

```

user@PE1> show multicast route extensive instance
vpna group 224.1.1.1
Family: INET

Group: 224.1.1.1
  Source: 192.168.1.2/32
  Upstream interface: fe-0/2/0.0
  Downstream interface list:
    mt-0/1/0.32768
  Session description: ST Multicast Groups
  Statistics: 7 kBps, 79 pps, 719738 packets
  Next-hop ID: 262144
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding

```

```
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
```

P-PIM and C-PIM on the Receiver PE Router

On the receiver PE router, multicast data packets received from the network are de-encapsulated as they are passed through the multicast tunnel de-encapsulation interface.

The P-PIM database on the receiver PE router contains two P-joins. One is for P-RP, and the other is for the sender PE router. For both entries, the OIL contains the multicast tunnel de-encapsulation interface from which the GRE header is stripped. The upstream interface for P-joins is the core-facing interface that faces towards the sender PE router.

Use the `show pim join extensive` command to verify that Router PE3 has the following state in its P-PIM database. The downstream neighbor interface points to the GRE de-encapsulation subinterface:

```
user@PE3> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.1.1.10
  Flags: sparse,rptree,wildcard
  Upstream interface: so-0/0/3.0
  Upstream neighbor: 10.12.100.21
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: mt-1/2/0.49152
      10.12.53.13 State: Join Flags: SRW Timeout: Infinity

Group: 239.1.1.1
  Source: 10.1.1.1
  Flags: sparse,spt
  Upstream interface: so-0/0/3.0
  Upstream neighbor: 10.12.100.21
  Upstream state: Join to Source
  Keepalive timeout: 351
  Downstream neighbors:
    Interface: mt-1/2/0.49152
      10.12.53.13 State: Join Flags: S Timeout: Infinity
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

On the VRF side of the receiver PE router, C-PIM installs a join entry in its C-PIM database. The OIL of this entry points to the local VPN interface, indicating active local receivers. The upstream protocol, interface, and neighbor of this entry point to the next-generation-MVPN network. Routers referenced in this topic are shown in ["Understanding Next-Generation MVPN Network Topology" on page 716](#).

Use the `show pim join extensive instance vpna 224.1.1.1` command to verify that Router PE3 has the following state in its C-PIM database:

```
user@PE3> show pim join extensive instance vpna
224.1.1.1
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Source: *
  RP: 10.12.53.1
  Flags: sparse,rptree,wildcard
  Upstream protocol: BGP
  Upstream interface: Through BGP
  Upstream neighbor: Through MVPN
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: so-0/2/0.0
      10.12.87.1 State: Join Flags: SRW Timeout: Infinity

Group: 224.1.1.1
  Source: 192.168.1.2
  Flags: sparse
  Upstream protocol: BGP
  Upstream interface: Through BGP
  Upstream neighbor: Through MVPN
  Upstream state: Join to Source
  Keepalive timeout:
  Downstream neighbors:
    Interface: so-0/2/0.0
      10.12.87.1 State: Join Flags: S Timeout: 195

Instance: PIM.vpna Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

The forwarding entry corresponding to the C-PIM entry on the receiver PE router uses the multicast tunnel de-encapsulation subinterface as the upstream interface.

Use the `show multicast route extensive instance vpna group 224.1.1.1` command to verify that Router PE3 has installed the following multicast forwarding entry for the local receiver:

```

user@PE3> show multicast route extensive instance
vpna group 224.1.1.1
Family: INET

Group: 224.1.1.1
  Source: 192.168.1.2/32
  Upstream interface: mt-1/2/0.49152
  Downstream interface list:
    so-0/2/0.0
  Session description: ST Multicast Groups
  Statistics: 1 kBps, 10 pps, 149 packets
  Next-hop ID: 262144
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

```

Provider Tunnels Signaled by RSVP-TE (Inclusive and Selective)

Junos OS supports signaling both inclusive and selective provider tunnels by RSVP-TE point-to-multipoint label-switched paths (LSPs). You can configure a combination of inclusive and selective provider tunnels per VPN.

- If you configure a VPN to use an inclusive provider tunnel, the sender PE router signals one point-to-multipoint LSP for the VPN.
- If you configure a VPN to use selective provider tunnels, the sender PE router signals a point-to-multipoint LSP for each selective tunnel configured.

Sender (ingress) PE routers and receiver (egress) PE routers play different roles in the point-to-multipoint LSP setup. Sender PE routers are mainly responsible for initiating the parent point-to-multipoint LSP and the sub-LSPs associated with it. Receiver PE routers are responsible for setting up state such that they can forward packets received over a sub-LSP to the correct VRF table (binding a provider tunnel to the VRF).

Inclusive Tunnels: Ingress PE Router Point-to-Multipoint LSP Setup

The point-to-multipoint LSP and associated sub-LSPs are signaled by the ingress PE router. The information about the point-to-multipoint LSP is advertised to egress PE routers in the PMSI attribute via BGP.

The ingress PE router signals point-to-multipoint sub-LSPs by originating point-to-multipoint RSVP path messages toward egress PE routers. The ingress PE router learns the identity of the egress PE routers from Type 1 routes installed in its `<routing-instance-name>.mvpn.0` table. Each RSVP path message carries an S2L_Sub_LSP object along with the point-to-multipoint session object. The S2L_Sub_LSP object carries a 4-byte sub-LSP destination (egress) IP address.

In Junos OS, sub-LSPs associated with a point-to-multipoint LSP can be signaled automatically by the system or via a static sub-LSP configuration. When they are automatically signaled, the system chooses a name for the point-to-multipoint LSP and each sub-LSP associated with it using the following naming convention.

Point-to-multipoint LSPs naming convention:

`<ingress PE rid>:<a per VRF unique number>:mvpn:<routing-instance-name>`

Sub-LSPs naming convention:

`<egress PE rid>:<ingress PE rid>:<a per VRF unique number>:mvpn:<routing-instance-name>`

Use the `show mpls lsp p2mp` command to verify that the following LSPs have been created by Router PE1:

Parent P2MP LSP: `10.1.1.1:65535:mvpn:vpna`

Sub-LSPs: `10.1.1.2:10.1.1.1:65535:mvpn:vpna` (Router PE1 to Router PE2) and

`10.1.1.3:10.1.1.1:65535:mvpn:vpna` (Router PE1 to Router PE3)

```

user@PE1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 2
To          From          State Rt P  ActivePath  LSPname
10.1.1.2    10.1.1.1      Up    0 *          10.1.1.2:10.1.1.1:65535:mvpn:vpna
10.1.1.3    10.1.1.1      Up    0 *          10.1.1.3:10.1.1.1:65535:mvpn:vpna
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
```

The values in this example are as follows:

- I-PMSI P2MP LSP name: 10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE2): 10.1.1.2:10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE3): 10.1.1.3:10.1.1.1:65535:mvpn:vpna

Inclusive Tunnels: Egress PE Router Point-to-Multipoint LSP Setup

An egress PE router responds to an RSVP path message by originating an RSVP reservation (RESV) message per normal RSVP procedures. The RESV message contains the MPLS label allocated by the egress PE router for this sub-LSP and is forwarded hop by hop toward the ingress PE router, thus setting up state on the network. Routers referenced in this topic are shown in "[Understanding Next-Generation MVPN Network Topology](#)" on page 716.

Use the `show rsvp session` command to verify that Router PE2 has assigned label 299840 for the sub-LSP 10.1.1.2:10.1.1.1:65535:mvpn:vpna:

```
user@PE2> show rsvp session
Total 0 displayed, Up 0, Down 0
Egress RSVP: 1 sessions
To          From      State   Rt Style  Labelin  Labelout  LSPname
10.1.1.2    10.1.1.1 Up 0 1 SE  299840   -         10.1.1.2:10.1.1.1:65535:mvpn:vpna
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Use the `show mpls lsp p2mp` command to verify that Router PE3 has assigned label 16 for the sub-LSP 10.1.1.3:10.1.1.1:65535:mvpn:vpna:

```
user@PE3> show mpls lsp p2mp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 1 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 1
```

```

To          From      State   Rt Style   Labelin  Labelout  LSPname
10.1.1.3    10.1.1.1 Up 0 1 SE    16        -
10.1.1.3:10.1.1.1:65535:mvpn:vpna
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Inclusive Tunnels: Egress PE Router Data Plane Setup

The egress PE router installs a forwarding entry in its `mpls` table for the label it allocated for the sub-LSP. The MPLS label is installed with a pop operation (a pop operation removes the top MPLS label), and the packet is passed on to the VRF table for a second route lookup. The second lookup on the egress PE router is necessary for the VPN multicast data packets to be processed inside the VRF table using normal C-PIM procedures.

Use the `show route table mpls label 16` command to verify that Router PE3 has installed the following label entry in its MPLS forwarding table:

```

user@PE3> show route table mpls label 16
+ = Active Route, - = Last Active, * = Both

16          *[VPN/0] 03:03:17
            to table vpna.inet.0, Pop

```

In Junos OS, VPN multicast routing entries are stored in the `<routing-instance-name>.inet.1` table, which is where the second route lookup occurs. In the example above, even though `vpna.inet.0` is listed as the routing table where the second lookup happens after the pop operation, internally the lookup is pointed to the `vpna.inet.1` table. Routers referenced in this topic are shown in ["Understanding Next-Generation MVPN Network Topology" on page 716](#).

Use the `show route table vpna.inet.1` command to verify that Router PE3 contains the following entry in its VPN multicast routing table:

```

user@PE3> show route table vpna.inet.1
vpna.inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

224.1.1.1,192.168.1.2/32*[MVPN/70] 00:04:10
                        Multicast (IPv4)

```


Use the `show multicast route extensive instance vpna` command to verify that Router PE3 contains the following VPN multicast forwarding entry corresponding to the multicast routing entry for the local join. The upstream interface points to `lsi.0` and the downstream interface (OIL) points to the `so-0/2/0.0` interface (toward local receivers). The `Upstream protocol` value is `MVPN` because the VPN multicast source is reachable via the next-generation MVPN network. The `lsi.0` interface is similar to the multicast tunnel interface used when PIM-based provider tunnels are used. The `lsi.0` interface is used for removing the top MPLS header.

```

user@PE3> show multicast route extensive instance
vpna
Family: INET

Group: 224.1.1.1
  Source: 192.168.1.2/32
  Upstream interface: lsi.0
  Downstream interface list:
    so-0/2/0.0
  Session description: ST Multicast Groups
  Statistics: 1 kBps, 10 pps, 3472 packets
  Next-hop ID: 262144
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

Family: INET6

```

The requirement for a double route lookup on the VPN packet header requires two additional configuration statements on the egress PE routers when provider tunnels are signaled by RSVP-TE.

First, since the top MPLS label used for the point-to-multipoint sub-LSP is actually tied to the VRF table on the egress PE routers, the penultimate-hop popping (PHP) operation is not used for next-generation MVPNs. Only ultimate-hop popping is used. PHP allows the penultimate router (router before the egress PE router) to remove the top MPLS label. PHP works well for VPN unicast data packets because they typically carry two MPLS labels: one for the VPN and one for the transport LSP.

After the LSP label is removed, unicast VPN packets still have a VPN label that can be used for determining the VPN to which the packets belong. VPN multicast data packets, on the other hand, carry only one MPLS label that is directly tied to the VPN. Therefore, the MPLS label carried by VPN multicast packets must be preserved until the packets reach the egress PE router. Normally, PHP must be disabled through manual configuration.

To simplify the configuration, PHP is disabled by default on Juniper Networks PE routers when you include the `mvpn` statement at the `[edit routing-instances routing-interface-name interface]` hierarchy level. PHP is also disabled by default when you include the `vrf-table-label` statement at the `[edit routing-instances routing-instance-name]` hierarchy level.

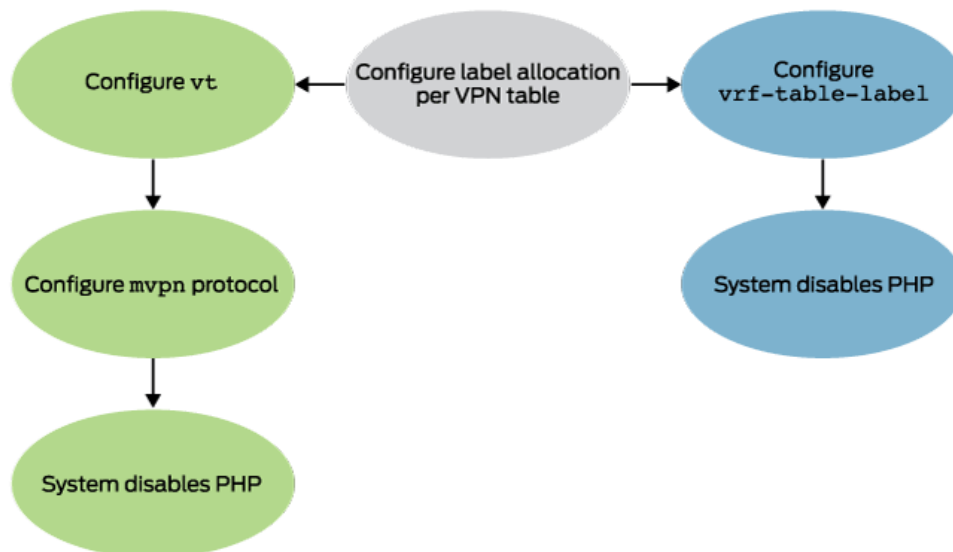
Second, in Junos OS, VPN labels associated with a VRF table can be allocated in two ways.

- Allocate a unique label for each VPN next hop (PE-CE interface). This is the default behavior.
- Allocate one label for the entire VRF table, which requires additional configuration. Only allocating a label for the entire VRF table allows a second lookup on the VPN packet's header. Therefore, PE routers supporting next-generation-MVPN services must be configured to allocate labels for the VRF table. There are two ways to do this as shown in [Figure 124 on page 1048](#).
 - One is by including a virtual tunnel interface named `vt` at the `[edit routing-instances routing-instance-name interfaces]` hierarchy level, which requires a Tunnel PIC.
 - The second is by including the `vrf-table-label` statement at the `[routing-instances routing-instance-name]` hierarchy level, which does not require a Tunnel PIC.

Both of these options enable an egress PE router to perform two route lookups. However, there are some differences in the way in which the second lookup is done

If the `vt` interface is used, the allocated label is installed in the `mpls` table with a `pop` operation and a forwarding next hop pointing to the `vt` interface.

Figure 124: Enabling Double Route Lookup on VPN Packet Headers



Use the `show route table mpls label 299840` command to verify that Router PE2 has installed the following entry and uses a `vt` interface in the `mpls` table. The label associated with the point-to-multipoint sub-LSP (299840) is installed with a `pop` and a forward operation with the `vt-0/1/0.0` interface being the next hop. VPN multicast packets received from the core exit the `vt-0/1/0.0` interface without their MPLS header, and the egress Router PE2 does a second lookup on the packet header in the `vpna.inet.1` table.

```
user@PE2> show route table mpls label 299840
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299840          *[VPN/0] 00:00:22
                > via vt-0/1/0.0, Pop
```

If the `vrf-table-label` is configured, the allocated label is installed in the `mpls` table with a `pop` operation, and the forwarding entry points to the `<routing-instance-name>.inet.0` table (which internally triggers the second lookup to be done in the `<routing-instance-name>.inet.1` table).

Use the `show route table mpls label 16` command to verify that Router PE3 has installed the following entry in its `mpls` table and uses the `vrf-table-label` statement:

```
user@PE3> show route table mpls label 16
mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

16              *[VPN/0] 03:03:17
                to table vpna.inet.0, Pop
```

Configuring label allocation for each VRF table affects both unicast VPN and MVPN routes. However, you can enable per-VRF label allocation for MVPN routes only if per-VRF allocation is configured via `vt`. This feature is configured via multicast and unicast keywords at the `[edit routing-instances <routing-instance-name> interface vt-x/y/z.0]` hierarchy level.

Note that including the `vrf-table-label` statement enables per-VRF label allocation for both unicast and MVPN routes and cannot be turned off for either type of routes (it is either on or off for both).

If a PE router is a bud router, meaning it has local receivers and also forwards MPLS packets received over a point-to-multipoint LSP downstream to other P and PE routers, then there is a difference in how the `vrf-table-label` and `vt` statements work. When, the `vrf-table-label` statement is included, the bud PE router receives two copies of the packet from the penultimate router: one to be forwarded to local receivers and the other to be forwarded to downstream P and PE routers. When the `vt` statement is included, the PE router receives a single copy of the packet.

Inclusive Tunnels: Ingress and Branch PE Router Data Plane Setup

On the ingress PE router, local VPN data packets are encapsulated with the MPLS label received from the network for sub-LSPs.

Use the `show rsvp session` command to verify that on the ingress Router PE1, VPN multicast data packets are encapsulated with MPLS label 300016 (advertised by Router P1 per normal RSVP RESV procedures) and forwarded toward Router P1 down the sub-LSPs `10.1.1.3:10.1.1.1:65535:mvpn:vpna` and `10.1.1.2:10.1.1.1:65535:mvpn:vpna`.

```

user@PE1> show rsvp session
Ingress RSVP: 2 sessions
To          From          State      Rt Style  Labelin  Labelout  LSPname
10.1.1.3    10.1.1.1    Up 0 1 SE   -        300016
10.1.1.3:10.1.1.1:65535:mvpn:vpna
10.1.1.2    10.1.1.1    Up 0 1 SE   -        300016
10.1.1.2:10.1.1.1:65535:mvpn:vpna
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

RFC 4875 describes a branch node as “an LSR that replicates the incoming data on to one or more outgoing interfaces.” On a branch Rrouter, the incoming data carrying an MPLS label is replicated onto one or more outgoing interfaces that can use different MPLS labels. Branch nodes keep track of incoming and outgoing labels associated with point-to-multipoint LSPs. Routers referenced in this topic are shown in ["Understanding Next-Generation MVPN Network Topology" on page 716](#).

Use the `show rsvp session` command to verify that branch node P1 has the incoming label 300016 and outgoing labels 16 for sub-LSP `10.1.1.3:10.1.1.1:65535:mvpn:vpna` (to Router PE3) and 299840 for sub-LSP `10.1.1.2:10.1.1.1:65535:mvpn:vpna` (to Router PE2).

```

user@P1> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

```

Transit RSVP: 2 sessions
To          From          State   Rt Style  Labelin  Labelout  LSPname
10.1.1.3 10.1.1.1 Up    0 1 SE    300016   16
10.1.1.3:10.1.1.1:65535:mvpn:vpna
10.1.1.2 10.1.1.1 Up    0 1 SE    300016  299840
10.1.1.2:10.1.1.1:65535:mvpn:vpna
Total 2 displayed, Up 2, Down 0

```

Use the `show route table mpls label 300016` command to verify that the corresponding forwarding entry on Router P1 shows that the packets coming in with one MPLS label (300016) are swapped with labels 16 and 299840 and forwarded out through their respective interfaces (so-0/0/3.0 and so-0/0/1.0 respectively toward Router PE2 and Router PE3).

```

user@P1> show route table mpls label 300016
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

300016          *[RSVP/7] 01:58:15, metric 1
                > via so-0/0/3.0, Swap 16
                via so-0/0/1.0, Swap 299840

```

Selective Tunnels: Type 3 S-PMSI Autodiscovery and Type 4 Leaf Autodiscovery Routes

Selective provider tunnels are configured by including the `selective` statement at the `[edit routing-instances routing-instance-name provider-tunnel]` hierarchy level. You can configure a threshold to trigger the signaling of a selective provider tunnel. Including the `selective` statement triggers the following events.

First, the ingress PE router originates a Type 3 S-PMSI autodiscovery route. The S-PMSI autodiscovery route contains the route distinguisher of the VPN where the tunnel is configured and the (C-S, C-G) pair that uses the selective provider tunnel.

In this section assume that Router PE1 is signaling a selective tunnel for (192.168.1.2, 224.1.1.1) and Router PE3 has an active receiver.

Use the `show route table vpna.mvpn.0 | find 3:` command to verify that Router PE1 has installed the following Type 3 route after the selective provider tunnel is configured:

```

user@PE1> show route table vpna.mvpn.0 | find
3:
3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1/240

```

```
*[MVPN/70] 00:05:07, metric2 1
  Indirect
```

Second, the ingress PE router attaches a PMSI attribute to a Type 3 route. This PMSI attribute is similar to the PMSI attribute advertised for inclusive provider tunnels with one difference: the PMSI attribute carried with Type 3 routes has its `Flags` bit set to `Leaf Information Required`. This means that the sender PE router is requesting receiver PE routers to send a Type 4 route if they have active receivers for the (C-S, C-G) carried in the Type 3 route. Also, remember that for each selective provider tunnel, a new point-to-multipoint and associated sub-LSPs are signaled. The PMSI attribute of a Type 3 route carries information about the new point-to-multipoint LSP.

Use the `show route advertising-protocol bgp 10.1.1.3 detail table vpna.mvpn | find 3:` command to verify that Router PE1 advertises the following Type 3 route and the PMSI attribute. The point-to-multipoint session object included in the PMSI attribute has a different port number (29499) than the one used for the inclusive tunnel (6574) indicating that this is a new point-to-multipoint tunnel.

```
user@PE1> show route advertising-protocol bgp
10.1.1.3 detail table vpna.mvpn | find 3:
* 3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1/240 (1 entry, 1 announced)
BGP group int type Internal
  Route Distinguisher: 10.1.1.1:1
  Nexthop: Self
  Flags: Nexthop Change
  Localpref: 100
  AS path: [65000] I
  Communities: target:10:1
  PMSI: Flags 1:RSVP-TE:label[0:0:0]:Session_13[10.1.1.1:0:29499:10.1.1.1]
```

Egress PE routers with active receivers should respond to a Type 3 route by originating a Type 4 leaf autodiscovery route. A leaf autodiscovery route contains a route key and the originating router's IP address fields. The `Route Key` field of the leaf autodiscovery route contains the original Type 3 route that is received. The originating router's IP address field is set to the router ID of the PE router originating the leaf autodiscovery route.

The ingress PE router adds each egress PE router that originated the leaf autodiscovery route as a leaf (destination of the sub-LSP for the selective point-to-multipoint LSP). Similarly, the egress PE router that originated the leaf autodiscovery route sets up forwarding state to start receiving data through the selective provider tunnel.

Egress PE routers advertise Type 4 routes with a route target that is specific to the PE router signaling the selective provider tunnel. This route target is in the form of `target:<rid of the sender PE>:0`. The sender PE router (the PE router signaling the selective provider tunnel) applies a special internal import

policy to Type 4 routes that looks for a route target with its own router ID. Routers referenced in this topic are shown in ["Understanding Next-Generation MVPN Network Topology" on page 716](#).

Use the `show route table vpna.mvpn | find 4:3:` command to verify that Router PE3 originates the following Type 4 route. The local Type 4 route is installed by the MVPN module.

```
user@PE3> show route table vpna.mvpn | find
4:3:
4:3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1:1.1.1.3/240
          *[MVPN/70] 00:15:29, metric2 1
          Indirect
```

Use the `show route advertising-protocol bgp 10.1.1.1 table vpna.mvpn detail | find 4:3:` command to verify that Router PE3 has advertised the local Type 4 route with the following route target community. This route target carries the IP address of the sender PE router (10.1.1.1) followed by a 0.

```
user@PE3> show route advertising-protocol bgp
10.1.1.1 table vpna.mvpn detail | find 4:3:
* 4:3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1:10.1.1.3/240 (1 entry, 1 announced)
BGP group int type Internal
  Nexthop: Self
  Flags: Nexthop Change
  Localpref: 100
  AS path: [65000] I
  Communities: target:10.1.1.1:0
```

Use the `show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__` command to verify that Router PE1 (the PE router signaling the selective provider tunnel) has applied the following import policy to Type 4 routes. The routes are accepted if their route target matches `target:10.1.1.1:0`.

```
user@PE1> show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__
Policy __vrf-mvpn-import-cmcast-leafAD-global-internal__:
Term unnamed:
from community __vrf-mvpn-community-rt_import-target-global-internal__ [target:10.1.1.1:0 ]
then accept
Term unnamed:
then reject
```

For each selective provider tunnel configured, a Type 3 route is advertised and a new point-to-multipoint LSP is signaled. Point-to-multipoint LSPs created by Junos OS for selective provider tunnels are named using the following naming conventions:

- Selective point-to-multipoint LSPs naming convention:

<ingress PE rid>:<a per VRF unique number>:mv<a unique number>:<routing-instance-name>

- Selective point-to-multipoint sub-LSP naming convention:

<egress PE rid>:<ingress PE rid>:<a per VRF unique>:mv<a unique number>:<routing-instance-name>

Use the `show mpls lsp p2mp` command to verify that Router PE1 signals point-to-multipoint LSP 10.1.1.1:65535:mv5:vpna with one sub-LSP 10.1.1.3:10.1.1.1:65535:mv5:vpna. The first point-to-multipoint LSP 10.1.1.1:65535:mvpn:vpna is the LSP created for the inclusive tunnel.

```

user@PE1> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 2
To      From      State      Rt P      ActivePath      LSPname
10.1.1.3 10.1.1.1 Up 0 *      10.1.1.3:10.1.1.1:65535:mvpn
:vpna
10.1.1.2 10.1.1.1 Up 0 *      10.1.1.2:10.1.1.1:65535:mvpn
:vpna
P2MP name: 10.1.1.1:65535:mv5:vpna, P2MP branch count: 1
To      From      State      Rt P      ActivePath      LSPname
10.1.1.3 10.1.1.1 Up 0 *      10.1.1.3:10.1.1.1:65535:mv5
:vpna
Total 3 displayed, Up 3, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

The values in this example are as follows.

- I-PMSI P2MP LSP name: 10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE2): 10.1.1.2:10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE3): 10.1.1.3:10.1.1.1:65535:mvpn:vpna
- S-PMSI P2MP LSP name: 10.1.1.1:65535:mv5:vpna
- S-PMSI P2MP sub-LSP name (to PE3): 10.1.1.3:10.1.1.1:65535:mv5:vpna

RELATED DOCUMENTATION

[Next-Generation MVPN Data Plane Overview | 730](#)

[Originating Type 1 Intra-AS Autodiscovery Routes Overview | 1033](#)

[Exchanging C-Multicast Routes | 1023](#)

Anti-spoofing support for MPLS labels in BGP/MPLS IP VPNs (Inter-AS Option B)

Service providers have traditionally adopted *Option A* VPN deployment scenarios instead of *Option B* because Option B is unable to ensure that the provider network is protected in the event of incorrect route distinguisher (RD) advertisements or spoofed MPLS labels.

Inter-AS Option B, however, can provide VPN services that are built using BGP based L3VPN. It is more scalable than the Option A alternative because Inter-autonomous system (AS) VPN routes are stored only in the BGP RIBs, as opposed to Option A which results in AS boundary routers (ASBRs) creating multiple VRF tables, each of which includes all IP routes.

Inter-AS Option B is also known as RFC 4364, *BGP/MPLS IP Virtual Private Networks*.

Junos OS Release 16.1 and later address the security shortcomings attributed to Option B. New features provide policy-based RD filtering (protection against MPLS label spoofing) to ensure that only RDs generated within the service provider domain are accepted. At the same time, the filtering can be used to filter loopback VPN-IPv4 addresses generated by PIM Rosen implementations from Cisco PEs, which can cause routing issues and traffic loss if imported into customer Virtual Routing and Forwarding (VRF) tables. These features are supported on M, MX, and T Series routers when using MPC1, MPC2, and MPC3D MPCs.

Inter-AS Option B uses BGP to signal VPN labels between ASBRs. The base MPLS tunnels are local to each AS, and stacked tunnels run from end-to-end between PE routers on the different AS VPN routes. The Junos OS anti-spoofing support for Option B implementations works by creating distinct MPLS forwarding table contexts. A separate `mpls.0` table is created for each set of VPN ASBR peers. As such, each MPLS forwarding table contains only the relevant labels advertised to the group of inter AS-Option B peers. Packets received with a different MPLS label are dropped. Option B peers are reachable through local interfaces that have been configured as part of the MFI (a new type of routing instance created for inter-AS BGP neighbors that require MPLS spoof-protection), so MPLS packets arriving from the Option B peers are resolved in the instance-specific MPLS forwarding table.

To enable anti-spoofing support for MPLS labels, configure separate instances of the new routing instance type, `mpls-forwarding`, on all MPLS-enabled Inter-AS links (which must be running a supported MPC). Then configure each Option B peer to use this routing instance as its `forwarding-context` under BGP. This forms the transport session with the peers and performs forwarding functions for traffic from

peers. Spoof checking occurs between any peers with different `mpls-forwarding` MFIs. For peers with the same `forwarding-context`, spoof-checking is not necessary because peers share the same `MFI.mpls.0` table.

Note that anti-spoofing support for MPLS labels is also supported on mixed networks, that is, those that include Juniper network devices that are not running a supported MPC, as long as the MPLS-enabled Inter-AS link is on a supported MPC. Any existing label-switched interface (LSI) features in the network, such as `vrf-table-label`, will continue to work as usual.

Inter-AS Option B supports graceful RE switchover (GRES), nonstop active routing (NSR), and in service software upgrades (unified ISSU).

RELATED DOCUMENTATION

[Solution to secure BGP Option B against MPLS label spoofing on MX Series routers](#)

instance-type

forwarding-context

BGP-MVPN SD-WAN Overlay

IN THIS SECTION

● [Spoke-side Configuration | 1057](#)

● [Hub-side Configuration | 1058](#)

In existing unicast SD-WAN deployments, the hub and spoke architecture is modified to ensure that the spokes are aware of all the routes. Multicast is overlaid on this existing deployment. No additional configuration is required.

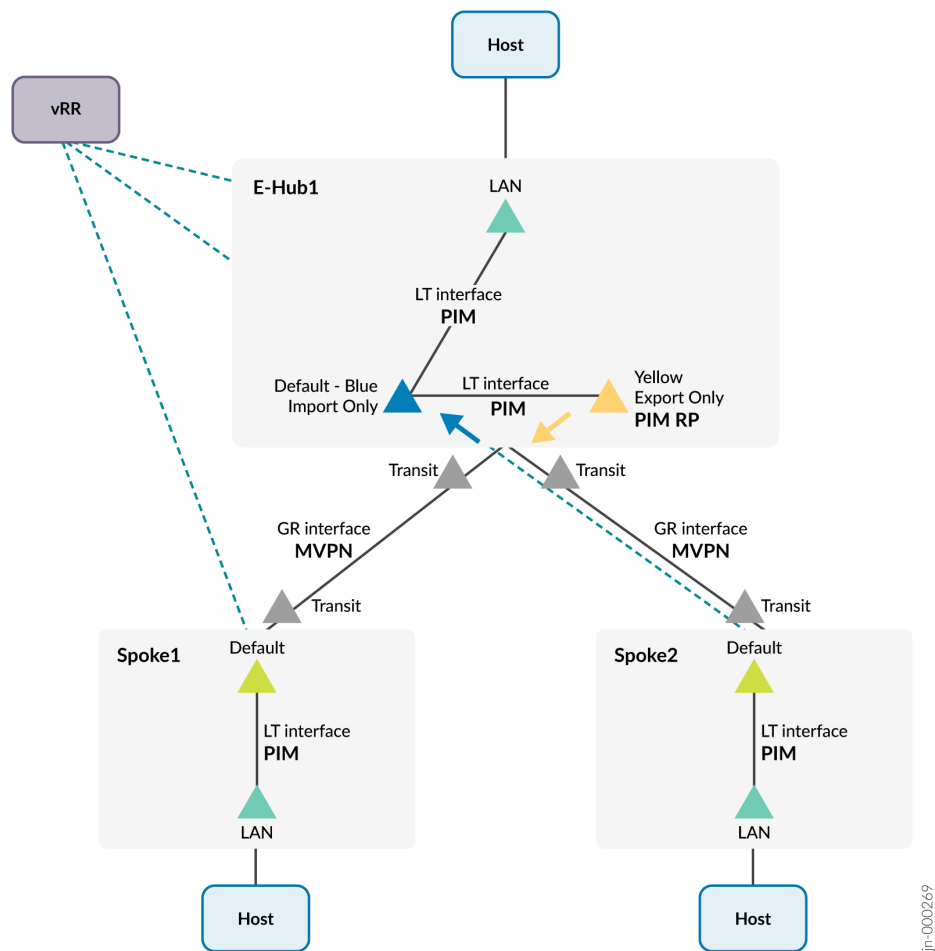
The unicast SD-WAN implementation introduces an additional VRF (yellow) on the hub that allows the traffic to traverse all routes. The yellow VRF only handles export of routes and the existing default VRF (blue) handles only import of routes. These VRFs are internally connected through the LT interface. This configuration allows the hub to advertise routes received from spokes to other spokes as if they originated from the hub itself. Therefore the spokes are aware of specific routes from other spokes.

Unlike unicast, a table jump next hop is not permissible in multicast and the traffic must traverse the extra hop through the LT interface wherever the reachability is through LT. Multicast does not apply advanced policy-based routing and hence will be applicable only in the default VRF instance.

In both the hub and the spoke, the VRF connected to the core can be considered as PE and the VRF connected to the host as CE. The multicast control plane uses the same path as that of unicast, using PIM over the LT interface and MVPN over the core GRE interface to exchange control plane information and update the forwarding routes.

Figure 125 on page 1057 shows the dual VRF configuration.

Figure 125: Enhanced hub and spoke model with Dual VRF



Spoke-side Configuration

On the spoke, PIM is enabled between the LAN VRF and the corresponding default VRF. The PIM joins traverse the LT interface from the LAN to the default VRF. MVPN is enabled on the default VRF. Source and RP reachability is configured through the LT interface so that the PIM joins are sent through the LT and the traffic follows the same path.

Hub-side Configuration

On the hub, PIM is enabled between the default blue and yellow VRFs. MVPN is also enabled on both, the default blue and yellow VRFs. There is no restriction on the import and export of MVPN routes, because it is exported and imported on both default blue and yellow VRFs. However, the yellow VRF will need to be configured as the MVPN sender site in SPT mode. The RP is configured on the export VRF (yellow) of the hub to ensure that type 5 routes are sent to all spokes to inform them of the active sources.

RELATED DOCUMENTATION

No Link Title

No Link Title

No Link Title

Configuring PIM Join Load Balancing

IN THIS CHAPTER

- [Use Case for PIM Join Load Balancing | 1059](#)
- [Configuring PIM Join Load Balancing | 1060](#)
- [PIM Join Load Balancing on Multipath MVPN Routes Overview | 1064](#)
- [Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN | 1068](#)
- [Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN | 1080](#)
- [Example: Configuring PIM Make-Before-Break Join Load Balancing | 1092](#)
- [Example: Configuring PIM State Limits | 1106](#)

Use Case for PIM Join Load Balancing

Large-scale service providers often have to meet the dynamic requirements of rapidly growing, worldwide virtual private network (VPN) markets. Service providers use the VPN infrastructure to deliver sophisticated services, such as video and voice conferencing, over highly secure, resilient networks. These services are usually loss-sensitive or delay-sensitive, and their data packets need to be delivered over a large-scale IP network in real time. The use of IP Multicast bandwidth-conserving technology has enabled service providers to exceed the most stringent service-level agreements (SLAs) and resiliency requirements.

IP multicast enables service providers to optimize network utilization while offering new revenue-generating value-added services, such as voice, video, and collaboration-based applications. IP multicast applications are becoming increasingly popular among enterprises, and as new applications start using multicast to deploy high-bandwidth and mission-critical services, it raises a new set of challenges for deploying IP multicast in the network.

IP multicast applications act as an essential communication protocol to effectively manage bandwidth and to reduce application server load by replicating the traffic on the network when the need arises. IP Protocol Independent Multicast (PIM) is the most important IP multicast routing protocol that is used to communicate between the multicast routers, and is the industry standard for building multicast distribution trees of receiving hosts. The multipath PIM join load-balancing feature in a multicast VPN

provides bandwidth efficiency by utilizing unequal paths toward a destination, improves scalability for large service providers, and minimizes service disruption.

The large-scale demands of service providers for IP access require Layer 3 VPN composite next hops along with external and internal BGP (EIBGP) VPN load balancing. The multipath PIM join load-balancing feature meets the large-scale requirements of enterprises by enabling `l3vpn-composite-nh` to be turned on along with EIBGP load balancing.

When the service provider network does not have the multipath PIM join load-balancing feature enabled on the provider edge (PE) routers, a hash-based algorithm is used to determine the best route to transmit multicast datagrams throughout the network. With hash-based join load balancing, adding new PE routers to the candidate upstream toward the destination results in PIM join messages being redistributed to new upstream paths. If the number of join messages is large, network performance is impacted because join messages are being sent to the new reverse path forwarding (RPF) neighbor and prune messages are being sent to the old RPF neighbor. In next-generation multicast virtual private network (MVPN), this results in multicast data messages being withdrawn from old upstream paths and advertised on new upstream paths, impacting network performance.

RELATED DOCUMENTATION

[PIM Join Load Balancing on Multipath MVPN Routes Overview | 1064](#)

Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN

Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN

Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs) and multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, the Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the `show pim join extensive` or `show pim source` commands.

```

user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
  Source: *
  RP: 10.255.245.6
  Flags: sparse,rptree,wildcard
  Upstream interface: t1-0/2/3.0
  Upstream neighbor: 192.168.38.57
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: t1-0/2/1.0
      192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
  Source: *

```

```

RP: 10.255.245.6
Flags: sparse,rptree,wildcard
Upstream interface: so-0/3/0.0
Upstream neighbor: 192.168.38.47
Upstream state: Join to RP
Downstream neighbors:
  Interface: t1-0/2/3.0
    192.168.38.16 State: JOIN Flags; SRW Timeout: 164

```

Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM sparse mode and join load balancing.

```

[edit protocols pim ]
user@host# set interface all mode sparse version 2
user@host# set join-load-balance

```

3. Then configure the static address of the RP.

```

[edit protocols pim rp]
user@host# set static address 10.10.10.1

```

4. Monitor the operation.

If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the `show pim interfaces` command.

```

user@host> show pim interfaces
Instance: PIM.master

Name           Stat Mode      IP V State NbrCnt JoinCnt  DR address
lo0.0          Up   Sparse       4 2 DR      0      0  10.255.168.58
pe-1/2/0.32769 Up   Sparse       4 2 P2P     0      0
so-0/3/0.0     Up   Sparse       4 2 P2P     1      1
t1-0/2/1.0     Up   Sparse       4 2 P2P     1      0
t1-0/2/3.0     Up   Sparse       4 2 P2P     1      1
lo0.0          Up   Sparse       6 2 DR      0      0  fe80::2a0:a5ff:4b7

```


Note that the two equal-cost paths shown by the `show pim interfaces` command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the `show pim neighbors detail` output:

```

user@host> show pim neighbors detail
Interface: so-0/3/0.0

Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1689116164
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
  BFD: Disabled
  Hello Option Holdtime: 105 seconds 102 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 792890329
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: t1-0/2/3.0

Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 678582286
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
  BFD: Disabled
  Hello Option Holdtime: 105 seconds 97 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1854475503
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the

interface or neighbor is restored, prior joins are not redistributed. The `clear pim join-distribution` command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

RELATED DOCUMENTATION

Load Balancing in Layer 3 VPNs

show pim interfaces

show pim neighbors

show pim source

clear pim join-distribution

PIM Join Load Balancing on Multipath MVPN Routes Overview

A multicast virtual private network (MVPN) is a technology to deploy the multicast service in an existing MPLS/BGP VPN.

The two main MVPN services are:

- Dual PIM MVPNs (also referred to as Draft-Rosen)
- Multiprotocol BGP-based MVPNs (also referred to as next-generation)

Next-generation MVPNs constitute the next evolution after the Draft-Rosen MVPN and provide a simpler solution for administrators who want to configure multicast over Layer 3 VPNs. A Draft-Rosen MVPN uses Protocol Independent Multicast (PIM) for customer multicast (C-multicast) signaling, and a next-generation MVPN uses BGP for C-multicast signaling.

Multipath routing in an MVPN is applied to make data forwarding more robust against network failures and to minimize shared backup capacities when resilience against network failures is required.

By default, PIM join messages are sent toward a source based on the reverse path forwarding (RPF) routing table check. If there is more than one equal-cost path toward the source [S, G] or rendezvous point (RP) [* , G], then one upstream interface is used to send the join messages. The upstream path can be:

- A single active external BGP (EBGP) path when both EBGP and internal BGP (IBGP) paths are present.
- A single active IBGP path when there is no EBGP path present.

With the introduction of the multipath PIM join load-balancing feature, customer PIM (C-PIM) join messages are load-balanced in the following ways:

- In the case of a Draft-Rosen MVPN, unequal EBGP and IBGP paths are utilized.
- In the case of next-generation MVPN:
 - Available IBGP paths are utilized when no EBGP path is present.
 - Available EBGP paths are utilized when both EBGP and IBGP paths are present.

This feature is applicable to IPv4 C-PIM join messages over the Layer 3 MVPN service.


By default, a customer source (C-S) or a customer RP (C-RP) is considered remote if the active **rt_entry** is a secondary route and the primary route is present in a different routing instance. Such determination is being done without taking into consideration the (C-*,G) or (C-S,G) state for which the check is being performed. The multipath PIM join load-balancing feature determines if a source (or RP) is remote by taking into account the associated (C-*,G) or (C-S,G) state.

When the provider network does not have provider edge (PE) routers with the multipath PIM join load-balancing feature enabled, hash-based join load balancing is used. Although the decision to configure this feature does not impact PIM or overall system performance, network performance can be affected temporarily, if the feature is not enabled.

With hash-based join load balancing, adding new PE routers to the candidate upstream toward the C-S or C-RP results in C-PIM join messages being redistributed to new upstream paths. If the number of join messages is large, network performance is impacted because of join messages being sent to the new RPF neighbor and prune messages being sent to the old RPF neighbor. In next-generation MVPN, this results in BGP C-multicast data messages being withdrawn from old upstream paths and advertised on new upstream paths, impacting network performance.

In [Figure 126 on page 1066](#), PE1 and PE2 are the upstream PE routers. Router PE1 learns route Source from EBGP and IBGP peers—the customer edge CE1 router and the PE2 router, respectively.

Figure 126: PIM Join Load Balancing



9040919

- If the PE routers run the Draft-Rosen MVPN, the PE1 router distributes C-PIM join messages between the EBGP path to the CE1 router and the IBGP path to the PE2 router. The join messages on the IBGP path are sent over a multicast tunnel interface through which the PE routers establish C-PIM adjacency with each other.

If a PE router loses one or all 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 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 new join messages get load-balanced across EBGp and IBGP paths, whereas the existing join messages on the multicast tunnel interface remain unaffected.

- If the PE routers run the next-generation MVPN, the PE1 router sends C-PIM join messages directly to the CE1 router over the EBGp path. There is no C-PIM adjacency between the PE1 and PE2 routers. Router PE3 distributes the C-PIM join messages between the two IBGP paths to PE1 and PE2. The Byte-wise-XOR hash algorithm is used to send the C-multicast data according to Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*.

Because the multipath PIM join load-balancing feature in a Draft-Rosen MVPN utilizes unequal EBGp and IBGP paths to the destination, loops can be created when forwarding unicast packets to the destination. To avoid or break such loops:

- Traffic arriving from a core or master instance should not be forwarded back to the core facing interfaces.
- A single multicast tunnel interface should either be selected as the upstream interface or the downstream interface.
- An upstream or downstream multicast tunnel interface should point to a non-multicast tunnel interface.

As a result of the loop avoidance mechanism, join messages arriving from an EBGp path get load-balanced across EIBGP paths as expected, whereas join messages from an IBGP path are constrained to choose the EBGp path only.

In [Figure 126 on page 1066](#), if the CE2 host sends unicast data traffic to the CE1 host, the PE1 router could send the multicast flow to the PE2 router over the MPLS core due to traffic load balancing. A data forwarding loop is prevented by ensuring that PE2 does not forward traffic back on the MPLS core because of the load-balancing algorithm.

In the case of C-PIM join messages, assuming that both the CE2 host and the CE3 host are interested in receiving traffic from the source (S, G), and if both PE1 and PE2 choose each other as the RPF neighbor toward the source, then a multicast tree cannot be formed completely. This feature implements mechanisms to prevent such join loops in the multicast control plane in a Draft-Rosen MVPN scenario.



NOTE: Disruption of multicast traffic or creation of join loops can occur, resulting in a multicast distribution tree (MDT) not being formed properly due to one of the following reasons:

- 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, *nonstop active routing* (NSR) is not supported and has no impact on the multicast traffic in a Draft-Rosen MVPN scenario.
- Any PE router in the provider network is running another vendor's implementation that does not apply the same hashing algorithm implemented in this feature.
- The multipath PIM join load-balancing feature has not been configured properly.

RELATED DOCUMENTATION

| *Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN*

Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN

IN THIS SECTION

- [Requirements | 1069](#)
- [Overview and Topology | 1069](#)
- [Configuration | 1073](#)
- [Verification | 1078](#)

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

This example requires the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, 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 127 on page 1073](#), 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 127 on page 1073](#), 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 127: PIM Join Load Balancing on Draft-Rosen MVPN



9040919

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1074](#)
- [Procedure | 1075](#)

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 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 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

```

Procedure

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 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 {
    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

IN THIS SECTION

- [Verifying PIM Join Load Balancing for Different Groups of Join Messages | 1078](#)

Confirm that the configuration is working properly.

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 EBGW 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

Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN

IN THIS SECTION

- [Requirements | 1080](#)
- [Overview and Topology | 1081](#)
- [Configuration | 1084](#)
- [Verification | 1090](#)

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

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 Bytewise-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 EBGP 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 128 on page 1084](#), 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 128 on page 1084](#), 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 128: PIM Join Load Balancing on Next-Generation MVPN



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1085](#)
- [Procedure | 1086](#)
- [Results | 1088](#)

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

```

Procedure

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

IN THIS SECTION

- [Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages | 1090](#)

Confirm that the configuration is working properly.

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](#)

Example: Configuring PIM Make-Before-Break Join Load Balancing

IN THIS SECTION

- [Understanding the PIM Automatic Make-Before-Break Join Load-Balancing Feature | 1092](#)
- [Example: Configuring PIM Make-Before-Break Join Load Balancing | 1093](#)

Understanding the PIM Automatic Make-Before-Break Join Load-Balancing Feature

The PIM automatic make-before-break (MBB) join load-balancing feature introduces redistribution of PIM joins on equal-cost multipath (ECMP) links, with minimal disruption of traffic, when an interface is added to an ECMP path.

The existing PIM join load-balancing feature enables distribution of joins across ECMP links. In case of a link failure, the joins are redistributed among the remaining ECMP links, and traffic is lost. The addition of an interface causes no change to this distribution of joins unless the `clear pim join-distribution` command is used to load-balance the existing joins to the new interface. If the PIM automatic MBB join load-balancing feature is configured, this process takes place automatically.

The feature can be enabled by using the `automatic` statement at the `[edit protocols pim join-load-balance]` hierarchy level. When a new neighbor is available, the time taken to create a path to the neighbor (standby path) can be configured by using the `standby-path-creation-delay seconds` statement at the `[edit protocols pim]` hierarchy level. In the absence of this statement, the standby path is created immediately, and the joins are redistributed as soon as the new neighbor is added to the network. For a join to be moved to the standby path in the absence of traffic, the `idle-standby-path-switchover-delay seconds` statement is configured at the `[edit protocols pim]` hierarchy level. In the absence of this statement, the join is not moved until traffic is received on the standby path.

```
protocols {
  pim {
    join-load-balance {
      automatic;
    }
    standby-path-creation-delay seconds;
  }
}
```

```
    idle-standby-path-switchover-delay seconds;  
  }  
}
```

Example: Configuring PIM Make-Before-Break Join Load Balancing

IN THIS SECTION

- [Requirements | 1093](#)
- [Overview | 1094](#)
- [Configuration | 1095](#)
- [Verification | 1101](#)

This example shows how to configure the PIM make-before-break (MBB) join load-balancing feature.

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers (M120 and M320 only), MX Series 5G Universal Routing Platforms, or T Series Core Routers (TX Matrix and TX Matrix Plus only).
- Junos OS Release 12.2 or later.

Before you configure the MBB feature, be sure you have:

- Configured the device interfaces.
- Configured an interior gateway protocol (IGP) for both IPv4 and IPv6 routes on the devices (for example, OSPF and OSPFv3).
- Configured multiple ECMP interfaces (logical tunnels) using VLANs on any two routers (for example, Routers R1 and R2).

Overview

IN THIS SECTION

- [Topology | 1094](#)


Junos OS provides a PIM automatic MBB join load-balancing feature to ensure that PIM joins are evenly redistributed to all upstream PIM neighbors on an equal-cost multipath (ECMP) path. When an interface is added to an ECMP path, MBB provides a switchover to an alternate path with minimal traffic disruption.

Topology

In this example, three routers are connected in a linear manner between source and receiver. An IGP protocol and PIM sparse mode are configured on all three routers. The source is connected to Router R0, and five interfaces are configured between Routers R1 and R2. The receiver is connected to Router R2, and PIM automatic MBB join load balancing is configured on Router R2.

[Figure 129 on page 1094](#) shows the topology used in this example.

Figure 129: Configuring PIM Automatic MBB Join Load Balancing



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1095](#)
- [Configuring PIM MBB Join Load Balancing | 1096](#)
- [Results | 1097](#)

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 (Source)

```
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim rp static address 10.255.12.34
set protocols pim rp static address abcd::10:255:12:34
```

Router R1 (RP)

```
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim rp local family inet address 10.255.12.34
set protocols pim rp local family inet6 address abcd::10:255:12:34
```

Router R2 (Receiver)

```
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim rp static address 10.255.12.34
set protocols pim rp static address abcd::10:255:12:34
set protocols mld interface ge-0/0/3 version 1
set protocols mld interface ge-0/0/3 static group ff05::e100:1 group-count 100
set protocols pim join load-balance automatic
```

```
set protocols pim standby-path-creation-delay 5
set protocols pim idle-standby-path-switchover-delay 10
```

Configuring PIM MBB Join Load Balancing

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 [Junos OS CLI User Guide](#).

To configure PIM MBB join load balancing across the setup:

1. Configure PIM sparse mode on all three routers.

```
[edit protocols pim interface all]
user@host# set mode sparse
user@host# set version 2
```

2. Configure Router R1 as the RP.

```
[edit protocols pim rp local]
user@R1# set family inet address 10.255.12.34
user@R1# set family inet6 address abcd::10:255:12:34
```

3. Configure the RP static address on non-RP routers (R0 and R2).

```
[edit protocols pim rp ]
user@host# set static address 10.255.12.34
user@host# set static address abcd::10:255:12:34
```

4. Configure the Multicast Listener Discovery (MLD) group for ECMP interfaces on Router R2.

```
[edit protocols mld interface ge-0/0/3]
user@R2# set version 1
user@R2# set static group ff05::e100:1 group-count 100
```

5. Configure the PIM MBB join load-balancing feature on the receiver router (Router R2).

```
[edit protocols pim]
user@R2# set join load-balance automatic
user@R2# set standby-path-creation-delay 5
user@R2# set idle-standby-path-switchover-delay 10
```

Results

From configuration mode, confirm your configuration by entering the `show protocols` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show protocols
ospf {
  area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/3.1;
    interface ge-0/0/3.2;
    interface ge-0/0/3.3;
    interface ge-0/0/3.4;
    interface ge-0/0/3.5;
  }
}
ospf3 {
  area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/3.1;
    interface ge-0/0/3.2;
    interface ge-0/0/3.3;
    interface ge-0/0/3.4;
    interface ge-0/0/3.5;
  }
}
pim {
  rp {
    static {
      address 10.255.12.34;
      address abcd::10:255:12:34;
    }
  }
}
```

```
interface all {
    mode sparse;
    version 2;
}
interface fxp0.0 {
    disable;
}
interface ge-0/0/3.1;
interface ge-0/0/3.2;
interface ge-0/0/3.3;
interface ge-0/0/3.4;
interface ge-0/0/3.5;
}
```

```
user@R1# show protocols
ospf {
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/3.1;
        interface ge-0/0/3.2;
        interface ge-0/0/3.3;
        interface ge-0/0/3.4;
        interface ge-0/0/3.5;
    }
}
ospf3 {
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/3.1;
        interface ge-0/0/3.2;
        interface ge-0/0/3.3;
        interface ge-0/0/3.4;
        interface ge-0/0/3.5;
    }
}
pim {
    rp {
        local {
            family inet {
                address 10.255.12.34;
            }
        }
    }
}
```

```
        family inet6 {
            address abcd::10:255:12:34;
        }
    }
}
interface all {
    mode sparse;
    version 2;
}
interface fxp0.0 {
    disable;
}
interface ge-0/0/3.1;
interface ge-0/0/3.2;
interface ge-0/0/3.3;
interface ge-0/0/3.4;
interface ge-0/0/3.5;
}
```

```
user@R2# show protocols
mld {
    interface ge-0/0/3.1 {
        version 1;
        static {
            group ff05::e100:1 {
                group-count 100;
            }
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-1/0/7.1;
        interface ge-1/0/7.2;
        interface ge-1/0/7.3;
        interface ge-1/0/7.4;
        interface ge-1/0/7.5;
        interface ge-0/0/3.1;
    }
}
ospf3 {
```

```
area 0.0.0.0 {
    interface lo0.0;
    interface ge-1/0/7.1;
    interface ge-1/0/7.2;
    interface ge-1/0/7.3;
    interface ge-1/0/7.4;
    interface ge-1/0/7.5;
    interface ge-0/0/3.1;
}
}
pim {
    rp {
        static {
            address 10.255.12.34;
            address abcd::10:255:12:34;
        }
    }
    interface all {
        mode sparse;
        version 2;
    }
    interface fxp0.0 {
        disable;
    }
    interface ge-1/0/7.1;
    interface ge-1/0/7.2;
    interface ge-1/0/7.3;
    interface ge-1/0/7.4;
    interface ge-1/0/7.5;
    interface ge-0/0/3.1;
    join-load-balance {
        automatic;
    }
    standby-path-creation-delay 5;
    idle-standby-path-switchover-delay 10;
}
```

Verification

IN THIS SECTION

- [Verifying Interface Configuration | 1101](#)
- [Verifying PIM | 1102](#)
- [Verifying the PIM Automatic MBB Join Load-Balancing Feature | 1104](#)

Verifying Interface Configuration

Purpose

Verify that the configured interfaces are functional.

Action

Send 100 (S,G) joins from the receiver to Router R2 . From the operational mode of Router R2, run the **show pim interfaces** command.

```
user@R2> show pim interfaces
```

```
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
```

Name	Stat	Mode	IP	V	State	NbrCnt	JoinCnt(sg/*g)	DR address
ge-0/0/3.1	Up	S	4	2	DR,NotCap	0	0/0	70.0.0.1
ge-1/0/7.1	Up	S	4	2	DR,NotCap	1	20/0	14.0.0.2
ge-1/0/7.2	Up	S	4	2	DR,NotCap	1	20/0	14.0.0.6
ge-1/0/7.3	Up	S	4	2	DR,NotCap	1	20/0	14.0.0.10
ge-1/0/7.4	Up	S	4	2	DR,NotCap	1	20/0	14.0.0.14
ge-1/0/7.5	Up	S	4	2	DR,NotCap	1	20/0	14.0.0.18

The output lists all the interfaces configured for use with the PIM protocol. The Stat field indicates the current status of the interface. The DR address field lists the configured IP addresses. All the interfaces are operational. If the output does not indicate that the interfaces are operational, reconfigure the interfaces before proceeding.

Meaning

All the configured interfaces are functional in the network.

Verifying PIM

Purpose

Verify that PIM is operational in the configured network.

Action

From operational mode, enter the **show pim statistics** command.

```
user@R2> show pim statistics
```

PIM Message type	Received	Sent	Rx errors
V2 Hello	4253	5269	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	1750	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
V2 State Refresh	0	0	0
V2 DF Election	0	0	0
V1 Query	0	0	0
V1 Register	0	0	0
V1 Register Stop	0	0	0
V1 Join Prune	0	0	0
V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

Global Statistics

Hello dropped on neighbor policy	0
Unknown type	0
V1 Unknown type	0
Unknown Version	0
Neighbor unknown	0
Bad Length	0
Bad Checksum	0
Bad Receive If	0
Rx Bad Data	0
Rx Intf disabled	0
Rx V1 Require V2	0
Rx V2 Require V1	0
Rx Register not RP	0
Rx Register no route	0
Rx Register no decap if	0
Null Register Timeout	0
RP Filtered Source	0
Rx Unknown Reg Stop	0
Rx Join/Prune no state	0
Rx Join/Prune on upstream if	0
Rx Join/Prune for invalid group	0
Rx Join/Prune messages dropped	0
Rx sparse join for dense group	0
Rx Graft/Graft Ack no state	0
Rx Graft on upstream if	0
Rx CRP not BSR	0
Rx BSR when BSR	0
Anycast Register Stop	0 0 0

The V2 Hello field lists the number of PIM hello messages sent and received. The V2 Join Prune field lists the number of join messages sent before the join-prune-timeout value is reached. If both values are nonzero, PIM is functional.

Meaning

PIM is operational in the network.

Verifying the PIM Automatic MBB Join Load-Balancing Feature

Purpose

Verify that the PIM automatic MBB join load-balancing feature works as configured.

Action

To see the effect of the MBB feature on Router R2:

1. Run the **show pim interfaces** operational mode command before disabling an interface.

```
user@R2> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
Name      Stat Mode IP V State      NbrCnt JoinCnt(sg/*g) DR address
ge-0/0/3.1 Up      S 4 2 DR,NotCap 0      0/0      70.0.0.1
ge-1/0/7.1 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.2
ge-1/0/7.2 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.6
ge-1/0/7.3 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.10
ge-1/0/7.4 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.14
ge-1/0/7.5 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.18
```

The JoinCnt(sg/*g) field shows that the 100 joins are equally distributed among the five interfaces.

2. Disable the ge-1/0/7.5 interface.

```
[edit]
user@R2# set interfaces ge-1/0/7.5 disable
user@R2# commit
```

3. Run the **show pim interfaces** command to check if load balancing of joins is taking place.

```
user@R2> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
Name      Stat Mode IP V State      NbrCnt JoinCnt(sg/*g) DR address
```

```

ge-0/0/3.1 Up      S 4 2 DR,NotCap 0      0/0      70.0.0.1
ge-1/0/7.1 Up      S 4 2 DR,NotCap 1      25/0     14.0.0.2
ge-1/0/7.2 Up      S 4 2 DR,NotCap 1      25/0     14.0.0.6
ge-1/0/7.3 Up      S 4 2 DR,NotCap 1      25/0     14.0.0.10
ge-1/0/7.4 Up      S 4 2 DR,NotCap 1      25/0     14.0.0.14

```

The JoinCnt(sg/*g) field shows that the 100 joins are equally redistributed among the four active interfaces.

4. Add the removed interface on Router R2.

```

[edit]
user@R2# delete interfaces ge-1/0/7.5 disable
user@R2# commit

```

5. Run the **show pim interfaces** command to check if load balancing of joins is taking place after enabling the inactive interface.

```

user@R2> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable
Name      Stat Mode IP V State      NbrCnt JoinCnt(sg/*g) DR address
ge-0/0/3.1 Up      S 4 2 DR,NotCap 0      0/0      70.0.0.1
ge-1/0/7.1 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.2
ge-1/0/7.2 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.6
ge-1/0/7.3 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.10
ge-1/0/7.4 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.14
ge-1/0/7.5 Up      S 4 2 DR,NotCap 1      20/0     14.0.0.18

```

The JoinCnt(sg/*g) field shows that the 100 joins are equally distributed among the five interfaces.



NOTE: This output should resemble the output in Step 1.

Meaning

The PIM automatic MBB join load-balancing feature works as configured.

SEE ALSO

[Configuring MLD | 59](#)

join-load-balance

Example: Configuring PIM State Limits

IN THIS SECTION

- [Controlling PIM Resources for Multicast VPNs Overview | 1106](#)
- [Example: Configuring PIM State Limits | 1109](#)

Controlling PIM Resources for Multicast VPNs Overview

IN THIS SECTION

- [System Log Messages for PIM Resources | 1108](#)

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 35 on page 1108](#) describes the different types of PIM system messages that you might see depending on your system log warning and log interval configurations.

Table 35: 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.

Table 35: PIM System Log Messages (Continued)

System Log Message	Definition
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.

Example: Configuring PIM State Limits**IN THIS SECTION**

- Requirements | 1110
- Overview | 1110
- Configuration | 1111
- Verification | 1122

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

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 130 on page 1111](#) shows the topology used in this example.

Figure 130: PIM State Limits Topology



"CLI Quick Configuration" on page 1111 shows the configuration for all of the devices in Figure 130 on page 1111. The section **Device PE1** below describes the steps for Device PE1.

Configuration

IN THIS SECTION

- Procedure | 1111

Procedure

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

```

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 register-limit family inet maximum 100
set routing-instances vpn-1 protocols pim 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

```

Device P

```

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
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

IN THIS SECTION

- [Monitoring the PIM State Information | 1122](#)

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.

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](#)



General Multicast Options

[Bit Index Explicit Replication \(BIER\) | 1125](#)

[Prevent Routing Loops with Reverse Path Forwarding | 1149](#)

[Use Multicast-Only Fast Reroute \(MoFRR\) to Minimize Packet Loss During Link Failures | 1173](#)

[Enable Multicast Between Layer 2 and Layer 3 Devices Using Snooping | 1232](#)

[Configure Multicast Routing Options | 1273](#)

[Controller-Based BGP Multicast Signaling | 1335](#)

Bit Index Explicit Replication (BIER)

IN THIS CHAPTER

- [Overview | 1125](#)
- [IS-IS extension for BIER | 1129](#)
- [Configuring BIER MVPN | 1132](#)
- [BIER Forwarding | 1136](#)

Overview

IN THIS SECTION

- [Benefits of BIER | 1126](#)
- [BIER Terminology | 1126](#)
- [BIER Architecture | 1128](#)

Bit Index Explicit Replication (BIER) architecture supports the optimal forwarding of multicast packets without requiring a legacy multicast protocol to build multicast trees or for intermediate routers to maintain any per-multicast-flow state which simplifies control and forwarding planes.

Multicast forwarding is achieved by encapsulating the multicast packet with a BIER header at the ingress router. The BIER header contains a bit string, with each bit representing an egress router in the multicast domain. This bit is mapped from a unique ID assigned to each BIER enabled egress or ingress router. A multicast flow overlay protocol like BGP-MVPN helps the BFERs (Bit Forwarding Egress Router) communicate with the BFIRs (Bit Forwarding Ingress Router) and tell them that they need to receive certain overlay multicast traffic, so that the BFIRs can set up an overlay multicast forwarding state with appropriate BIER encapsulation information.

Benefits of BIER

BIER implementation results in considerable simplification of the multicast network, due to the elimination of:

- the per-flow state.
- explicit tree-building protocols.

BIER Terminology

The following terms are widely used in this topic:

Table 36: BIER terminology

Term	Definition
BIER	Bit Index Explicit Replication.
BFR	A Bit Forwarding Router is a BIER enabled router with a unique BFR prefix and optionally, a BFR-ID . A BFR establishes router adjacencies, computes the BIER routing and forwarding tables, and forwards or replicates BIER packets. There are three types of BFRs - BFIR, BFER and transit BFR.
BFIR	A Bit Forwarding Ingress Router is the first router in the BIER domain entered by a multicast packet. The BFIR adds a BIER header and forwards the packet using the BIER forwarding table.
BFER	A Bit Forwarding Egress Router is the last router that processes a BIER packet in a BIER domain. The BFER removes the BIER header before forwarding the packet.
Transit BFR	A transit Bit Forwarding Router is a router in the BIER domain that receives a multicast data packet from a BFR in the same BIER domain, and forwards the packet to another BFR within the same BIER domain.

Table 36: BIER terminology (Continued)

Term	Definition
BFR prefix	A BFR prefix is typically the configured loopback address of a BFR and must be a routable IP address in the BIER domain.
BFR-ID	A BFR-ID is a number in the range of 1-65535. When a BFR-ID is encoded in a packet, it is converted to a Set-ID (SI) and a bit in the bit string (BSL). Any BFR that is an ingress or egress router in a BIER domain is assigned a BFR-ID. A BFR-ID must be unique within that BIER sub-domain. Each of these BFR-IDs is encoded as a bit in the bit string that is carried in the BIER header with each multicast packet.
bit string	A bit string is a part of the BIER header, with each bit representing a BFER in the multicast domain.
BSL	The length of the bit string. By default, Junos only supports 256.
SD	Sub-domain. A BIER domain is a connected set of BFRs, each with a unique BFR-ID. A BIER domain can be divided into multiple sub-domains for various reasons. For example, it could represent a specific topology within a BIER domain.
Set-ID	In cases where the number of BFIRs/BFERs in the network is greater than the bitstringlength (BSL), they must be divided into multiple sets. Each set is identified by a Set-ID. However, this will require a packet to be replicated to each set. Junos currently supports 1 to 4 sets.
F-BM	Forwarding Bit Mask is the property of a BFR-neighbor which represents all the BFERs that are reachable through that BFR neighbour.

Table 36: BIER terminology (*Continued*)

Term	Definition
BIFT	<p>The BIER Forwarding Table is used by the BFR to identify which neighbors the packet should be sent to. The BIFT is created by calculating how each BFER is reached in the IGP routing underlay. Each BIFT entry maps a BFER's BFR-ID % BSL to a BFR neighbor and its corresponding F-BM. Each BIFT is specific to a particular sub-domain and Set-Identifier (SI).</p> <p>A BIFT's name is in the form of :bier-<subdomain-id>-<set-id>.bier.0.</p>
BIRT	<p>The BIER Routing Table.</p> <p>The routing protocol underlay (IS-IS for example) router advertises the BFR-prefix and sub-domain information inside IS-IS sub-TLVs which are flooded throughout the IS-IS domain. On the receiver side, IS-IS routers parse this BIER sub-domain information associated with the BFR-Prefix and derive the route and next-hop information. These routes are then installed into a routing table called BIRT.</p> <p>Routes in BIRT are keyed in on BIER prefixes. A BIRT's name is in the form of :bier-<subdomain-id>.inet.9.</p>

BIER Architecture

BIER architecture has three layers:

- **Routing underlay** – The underlay establishes adjacencies between pairs of BFRs, and determines one or more optimal paths from a BFR to a set of target BFRs.
- **BIER layer** - The BIER layer consists of the protocol and procedures that are used in order to transmit a multicast data packet across a BIER domain, from the BFIR to the BFERs.
- **Multicast flow overlay** - The overlay consists of a set of protocols and procedures that enable the following set of functions.

- When a BFIR receives a multicast data packet from outside the BIER domain, the BFIR must determine the set of BFERs for that packet. This information is provided by the multicast flow overlay.
- When a BFER receives a BIER-encapsulated packet from inside the BIER domain, the BFER must determine how to forward the packet. This information is provided by the multicast flow overlay. BGP-MVPN is an example of a multicast flow overlay.

IS-IS extension for BIER

IN THIS SECTION

- [Advertising BIER Info sub-TLV in IS-IS | 1129](#)
- [Advertising BIER MPLS Encapsulation sub-sub-TLV in IS-IS | 1130](#)
- [Advertising BIER-Prefix in IS-IS | 1131](#)
- [Installing BFR-Prefix to the BIER Routing Table | 1131](#)
- [BIER protocol and IS-IS interactions | 1132](#)

The IS-IS extension for BIER is defined via RFC 8401.

Junos OS Evolved supports the advertisement of BIER information of one or more BIER sub-domains using IS-IS as the IGP underlay. Key BIER information like BFR-IDs and BFR-Prefixes in each sub-domain are flooded through the IS-IS domain to generate the BIER forwarding table.

Advertising BIER Info sub-TLV in IS-IS


The BIER info sub-TLV carries BIER sub-domain information. This sub-TLV advertises a single <MT,SD> (multi-topology/sub-domain) combination followed by optional sub-sub-TLVs as described in the next section. Junos OS Evolved supports IPv4-unicast and IPv6-unicast IS-IS topologies.



NOTE: Multi-topology scenarios are not supported.

The figure below illustrates the BIER info sub-TLV.

Figure 131: BIER sub-TLV



IS-IS carries BIER info sub-TLV within IS-IS TLVs 135, 236 or 237.


- IS-IS advertises BIER info sub-TLV in TLV 135 for BIER sub-domains with BFER-Prefix (IPv4) and IS-IS default topology.
- IS-IS advertises BIER info sub-TLV in TLV 236 for BIER sub-domains with BFER-Prefix (IPv6) and IS-IS default topology.
- IS-IS advertises BIER info sub-TLV in TLV 237 for BIER sub-domains with BFER-Prefix (IPv6) and IS-IS IPv6 unicast topology.

Advertising BIER MPLS Encapsulation sub-sub-TLV in IS-IS

This sub-sub-TLV carries the information for the BIER MPLS encapsulation including the label range for a specific BitStringLength for a certain <MT,SD> tuple. The sub-sub-TLV is advertised within the BIER info sub-TLV and may appear multiple times within a single BIER info sub-TLV.

The figure below illustrates the BIER info sub-TLV.

Figure 132:



The Max SI value is the set-identifier and is configured under `number-sets`. For example in the snippet below:

```
sub-domain 32 {
    encapsulation mpls bitmask-length 256 number-sets 4
}
```

The `number-sets` is set to 4. The range is 1-16. Based on this value BIER assigns four contiguous labels from the dynamic pool, for example, labels L1 to L4. L1 belongs to SI=0, L2 to S1=1 and so on. With the above configuration example, IS-IS sub-TLV advertises with Max SI=4 and label as L1.

Advertising BIER-Prefix in IS-IS

BIER sub-domain related configuration parameters are set under `protocols bier`.

To include the BIER sub-domain in the IS-IS sub-TLVs, configure them under `protocols isis`:

```
protocols isis {
    bier-sub-domain 32
}
```

A BFR's BFR-Prefix is an IP address (IPv4 or IPv6) which is generally the loopback IP address of the BIER router. Typically, IS-IS advertises the BFR-Prefix in one of the TLVs 135/236/237 depending on the IS-IS topology.

When `bier-sub-domain` is configured under `protocols isis`, the information of the BIER sub-domain (BIER info sub-TLV + BIER MPLS Encapsulation sub-sub-TLV) is advertised inside the IS-IS prefix TLV to extend BIER functionality. Multiple sub-domains can advertise with the same loopback address.



NOTE: IS-IS TLVs have a maximum length of 255 bytes. When the maximum number of sub-domains that can be supported is exceeded, another loopback IP address must be configured.

Installing BFR-Prefix to the BIER Routing Table

IS-IS routers advertise the BFR-Prefix and sub-domain information using the sub-TLV. IS-IS floods this BIER information throughout the IS-IS domain. This allows the IS-IS nodes build the prefix and associated next-hop information to install the BIER route to the routing information base (RIB) table. On the receiver side, IS-IS nodes parse this BIER sub-domain information associated with the BFR-Prefix

and derive the route and next-hop information. IS-IS then installs this BIER route in the BIER routing table (BIRT).



NOTE: For this implementation, it is assumed that *a//IS-IS* nodes are BIER enabled.

BIER protocol and IS-IS interactions

IS-IS downloads the BIER route to the BIER routing table (BIRT). BIRT routes do not go directly to the forwarding table. These routes are consumed by the BIER module by listening to the route's updates from the BIRT table. The module then derives the new route and next-hop information from the BIRT route and eventually downloads it to the PFE for data forwarding.

The BIRT route table name is in the format: `bier-<subdomain-id>-<set-id>.bier.0`.

RELATED DOCUMENTATION

bier

bier-sub-domain (Protocols IS-IS)

Configuring BIER MVPN

IN THIS SECTION

- [Configuring BIER Provider Tunnels | 1133](#)
- [Verifying BIER in MVPN | 1133](#)

A Bit Forwarding Ingress Router (BFIR) encapsulates the incoming non-BIER multicast packets with the BIER header. For it to know which bit string to use, a multicast flow overlay protocol is needed so that the BFERs can tell BFIRs that the BFERs need to receive certain overlay (for example IP) multicast traffic. This way BFIRs can set up an overlay multicast forwarding state with appropriate BIER encapsulation information. BGP-MVPN is one such multicast overlay protocol, as it already has mechanisms for egress PEs (BFERs) to notify ingress PEs (BFIRs) that the BFERs need to receive traffic for certain (C-S/*, C-G/*).

Configuring BIER Provider Tunnels

You must include the `bier` statement at the `[edit protocols]` hierarchy level to enable BIER on the router.

To configure a Bit Index Explicit Replication (BIER) provider tunnel for a multicast VPN, include the `bier` statement:

```
bier {
  label label;
  subdomain-id subdomain-id;
}
```

You can include this statement at the following hierarchy level:

```
[edit routing-instances routing-instance-name provider-tunnel]
```

You can also configure `bier` for selective provider tunnels by configuring the following statement:

```
user@router# set routing-instances routing-instance-name provider-tunnel selective group
multicast-prefix/prefix-length source ip-prefix/prefix-length bier
```

The provider tunnel label should be the same as the configured static `vrf-table-label` so that this label can be assigned to the `lsi` interface associated with that routing-instance.

```
user@router# set routing-instances routing-instance-name vrf-table-label static label
```

In order to avoid interoperability issues with other vendors, the PMSI length should be the same size as the BIER prefix. It can be either IPv4 or IPv6.

Verifying BIER in MVPN

Issue the `show mvpn instance` command to display provider tunnel information.

```
user@router> show mvpn instance

Instance : vrf1
MVPN Mode : SPT-ONLY
Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
Hot Root Standby: Disabled. Reason: Not enabled by configuration.
```

```

Provider tunnel: I-P-tnl:BIER: subdomain-id 10 bfr-id 1 label 990001
Neighbor                Inclusive Provider Tunnel                Label-In
St      Segment
10.2.2.2                BIER: subdomain-id 10 bfr-id 2 bier-prefix 10.2.2.2    990001
10.3.3.3                BIER: subdomain-id 10 bfr-id 3 bier-prefix 10.3.3.3    990001
10.4.4.4

C-mcast IPv4 (S:G)      Provider Tunnel                Label-In
St  FwdNh  Segment
172.11.21.21/32:232.252.1.1/32  BIER: subdomain-id 10 bfr-id 1 label
990001                RM    M-8169
172.11.21.21/32:232.252.1.2/32  BIER: subdomain-id 10 bfr-id 1 label
990001                RM    M-8169
172.11.21.21/32:232.252.2.1/32  BIER: subdomain-id 10 bfr-id 1 label
990001                RM    M-8169
172.11.21.21/32:232.252.2.2/32  BIER: subdomain-id 10 bfr-id 1 label
990001                RM    M-8169

```

From the output above, it can be seen that two neighbors advertise the BIER tunnel in sub-domain 10 in their I-PMSI A-D routes, with BFR-ID 2 and 3 respectively. This router also advertises a BIER tunnel in sub-domain 10 with its own BFR-ID (1) and label 999901 to identify this VPN.

The I-PMSI tunnel is used for four (S,G) flows and the forwarding next hop used for those flows is a next hop with ID 8169, displayed as M-8169.

Issue the `show multicast route extensive instance instance-name` command to review those flows in detail.

```

user@router> show multicast route extensive instance vrf1
Instance: vrf1 Family: INET

Group: 232.252.1.1
  Source: 172.11.21.21/32
  Upstream interface: ge-0/0/0.1
  Downstream interface list:
    Push 999901 bier bitstring
00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000006: label 800000
  Number of outgoing interfaces: 0
  Session description: Source specific multicast
  Statistics: 0 kBps, 0 pps, 0 packets
  Next-hop ID: 8169
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding

```



```
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
Uptime: 01:26:43
.....trimmed
```

Push 999901 - this label identifies the VPN.

bier bitstring - the bit string to be encoded in the BIER header. In this example, the 2nd and 3rd bits are set, indicating that these two corresponding BFERs are to receive traffic.

label 800000 - the label that this router advertised for the corresponding BIFT. When the (S,G) packet is received on PE-CE interface, it matches this route so the VPN label is imposed. The BIER header is encoded with the bit string and this label, and then the packet is treated as if the BIER packet was just received on a core interface.

To display the local tunnel name, issue the `show multicast route extensive instance instance-name display-tunnel-name` command.

```
user@router> show multicast route extensive instance vrf1 display-tunnel-name
Instance: vrf1 Family: INET

Group: 232.252.1.1
  Source: 172.11.21.21/32
  Upstream interface: ge-0/0/0.1
  Downstream interface list:
    bier:mvpn:3
  Number of outgoing interfaces: 0
  Session description: Source specific multicast
  Statistics: 0 kBps, 0 pps, 0 packets
  Next-hop ID: 8169
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 01:26:43
```

RELATED DOCUMENTATION

bier

bier-sub-domain (Protocols IS-IS)

BIER Forwarding

IN THIS SECTION

- [Overview | 1136](#)
- [BIER forwarding in MVPN | 1139](#)

Overview

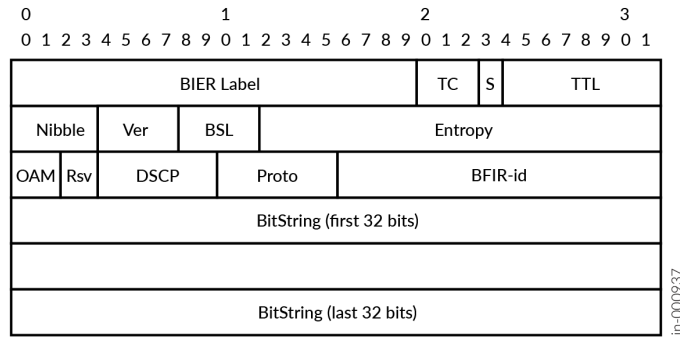
BIER forwarding is based on a bit string in the BIER header. Each bit set in the bit string represents a BFER (BIER Forwarding Egress Router) that should receive the multicast packet. The bit is mapped from a unique BFR-ID assigned to each BFER or BFIR. Note that a transit BIER Forwarding Router (BFR) that is neither a BFIR nor a BFER is not assigned a BFR-ID.

A BFR looks up the BIER Forwarding Table (BIFT) using the index of a set bit in the bit string of an incoming BIER packet. Since each bit set in the bit string corresponds to a specific BFR, the forwarding BFR is able to determine the neighbor the packet should be forwarded to.

The BIFT is created by calculating how each BFER is reached in the IGP routing underlay, e.g., a topology. A topology may have multiple BIER sub-domains, though a sub-domain is associated with only one topology. A BFR may belong to multiple sub-domains, and have different BFR-IDs for different sub-domains.

When the number of BFERs is greater than the BSL, multiple copies are sent, one for each set of BFERs. Each copy (for a different set of BFERs) uses a different BIFT for forwarding because the same bit in different copies is for different BFERs (e.g. bit-1 in copy-1 is for BFER 1 but in copy-2 it's for BFER 257, considering the BitStringLength is 256 bits).

An MPLS label is inserted at the start of the BIER header within an MPLS data plane to indicate that a BIER header is about to follow.



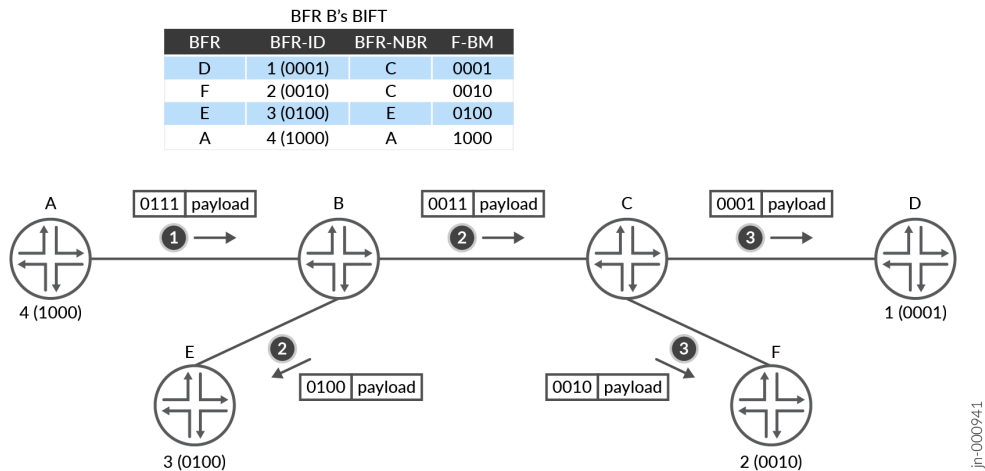
This BIER label also indicates which BIFT is to be used to forward the packets.

Routing protocol extensions are used to signal encapsulation information like BIER labels and calculate the BIFTs for all <sub-domain, BitStringLength, set> tuples that a BFR supports. The BIER information is attached to route advertisements of BFR prefixes i.e., loopback addresses.

When the route to a BFER prefix is calculated for a <sub-domain, BitStringLength> tuple, an entry is added to a corresponding BIFT with the key being (BFR-ID % BitStringLength) and the next hop information being <next hop neighbor, neighbor's BIER label for the <sub-domain, BitStringLength, set>, neighbor's F-BM>. The neighbor's F-BM is the logical OR of bits of all BFERs in this set reached by this neighbor.

Let's take an example of a simple network topology that consists of six BFRs - A, B, C, D, E, and F to illustrate this concept.

Figure 133: Bit forwarding mechanism



Out of these, routers A, D, E, and F act as edge routers i.e. either BFER or BFIR. Each has a unique BFR-ID assigned: 1 for D, 2 for F, 3 for E, and 4 for A. The BSL (bitstringlength) is set to 4 bits, and the SI (Set-ID) is currently 0. All BFRs belong to sub-domain 32.

Each of these BFR-IDs are mapped to a bit in a 4 bit bit string. This translates to the following assignments:

Router D: 0:0001

Router F: 0:0010

Router E: 0:0100

Router A: 0:1000

The remaining BFRs (B and C) are transit BFRs and are not assigned BFR-IDs.

When BFR A sends a BIER packet to BFR B with bit string 0111, BFR B determines that the packet is meant for BFRs E, F, and D. BFR B then sends two copies of the same packet, one to BFR E and another to BFR C. The copy to BFR C is sent with an updated bit string 0011, which corresponds to BFRs D and F. When the packet reaches BFR C, it decides to send one copy of the packet to BFR D and another to BFR F.

This forwarding decision is based on the BIFT of BFR B.

Observe that the BIFT lists the BFR neighbors through which BFRs A, D, E, and F can be reached. The F-BM is the property of the BFR neighbor, which represents all the BFRs that are reachable through that neighbor. For instance, BFR neighbor C's F-BM is 0011 which means that BFR C can cover for BFRs D and F. Router D has a F-BM of 0001 and BFR F has a F-BM of 0010. The F-BM of BFR C is the logical 'OR' of bits of all BFRs in this set reached by this neighbor which is 0011.

BIFT Route Lookup


Using the same example from above, when BFR A sends a BIER packet to BFR B with bit string 0111, BFR B uses the index of the lowest bit that is set in the incoming bit string (the right-most bit) to lookup the BIFT. The first entry (index one) in the routing table matches the lookup which implies that the BIER packet should be forwarded to BFR Neighbor C that has an F-BM of 0011. BFR B sends the copy to BFR C with the updated bit string 0011, which is the logical 'AND' of the incoming bit string and BFR C's F-BM.

Next, BFR B clears some bits in the incoming bit string using BFR C's F-BM, and uses the index of lowest set bit in the resulting bit string to lookup the BIFT. The table lookup results in BFR-neighbor E with F-BM 0100. BFR B sends a copy to BFR E with the updated bit string 0100 which is a logical 'AND' operation of the incoming bit string and BFR-neighbor E's F-BM.

This process continues until the incoming bit string becomes all zero.

BIER forwarding in MVPN

Lets take an example to illustrate the flow of packets through MVPN tunnels.



On PE1 (ingress):

Issue the show mvpn instance vrf1 command to display provider tunnel vrf1 information.

```
user@routerPE1> show mvpn instance vrf1
```

```
MVPN instance:
```

```
Instance : vrf1
```

```
MVPN Mode : SPT-ONLY
```

```
Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
```

```
Hot Root Standby: Disabled. Reason: Not enabled by configuration.
```

```
Provider tunnel: I-P-tnl:BIER: subdomain-id 10 bfr-id 1 label 990001
```

Neighbor	Inclusive Provider Tunnel	Label-In
St Segment		
10.2.2.2	BIER: subdomain-id 10 bfr-id 2 bier-prefix 10.2.2.2	990001
10.3.3.3	BIER: subdomain-id 10 bfr-id 3 bier-prefix 10.3.3.3	990001
10.4.4.4		

C-mcast IPv4 (S:G)	Provider Tunnel	Label-In
St FwdNh Segment		
172.11.21.21/32:232.252.1.1/32	BIER: subdomain-id 10 bfr-id 1 label	
990001	RM M-8169	
172.11.21.21/32:232.252.1.2/32	BIER: subdomain-id 10 bfr-id 1 label	
990001	RM M-8169	
172.11.21.21/32:232.252.2.1/32	BIER: subdomain-id 10 bfr-id 1 label	
990001	RM M-8169	

```
172.11.21.21/32:232.252.2.2/32  BIER: subdomain-id 10 bfr-id 1 label
990001                          RM    M-8169
```



NOTE: Issue the `show multicast route extensive instance vrf1 display-tunnel-name` and `show mvpn instance vrf1 display-tunnel-name` to view the BIER tunnel's name.

The inclusive tunnel has two MVPN neighbors, 10.2.2.2 with BFR-ID 2 and 10.3.3.3 with BFR-ID 3. Therefore bits 2 and 3 are set. Neighbor 10.4.4.4 did not advertise any tunnels so it is not considered.

The label 990001 is configured statically by issuing the `set routing-instances routing-instance-name vrf-table-label static label` statement.

In the FBM, 8 values are present for a bitstringlength of 256 bytes, separated by “:”, with each value representing 32 BFIR/BFERS.

Issue the `show multicast route instance vrf1 extensive` command to view multicast route information.

```
user@routerPE1> show multicast route instance vrf1 extensive
Instance: vrf1 Family: INET

Group: 232.252.1.1
  Source: 172.11.21.21/32
  Upstream interface: et-0/0/10.1
  Downstream interface list:
    Push 990001, BS:0:0:0:0:0:0:0:0:00000006, label 16
  Number of outgoing interfaces: 1
  Session description: Source specific multicast
  Statistics: 0 kbps, 1 pps, 44 packets
  Next-hop ID: 8169
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 01:26:43
.....trimmed
```

A (172.11.21.21, 232.252.1.1) packet arriving in the VRF matches the above route. Label 990001 is imposed first, so that the egress PEs can find the corresponding VRF when they remove the BIER header. Then a BIER header is imposed with bit string 0:0:0:0:0:0:0:0:00000006 and BIER label 16, which is advertised by this BFIR for the sub-domain of the provider tunnel. The result is then treated as if it were a BIER packet received from another BFR.

Label 16 is looked up in the local mpls.0 table. Issue the show route table mpls.0 protocol bier label 16 extensive command to view the BIFT that is looked up.

```

user@routerPE1> show route table mpls.0 protocol bier label 16 extensive

mpls.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)
16 (1 entry, 1 announced)
TSI:
KRT in-kernel 16      /52 -> {[6025]}
Opaque data client: BIER
Opaque data: TLV type:32820 APP flag:0x80
Address: 0x56026719ae20
Opaque-data reference count: 24
Stats ID Group: Kernel ID = 4097, Stats IDs = { 4026531842 }
    *BIER Preference: 70
        Next hop type: Multicast (IPv4) Composite, Next hop index: 6025
        Address: 0x560263988dc4
        Next-hop reference count: 2, key opaque handle: 0x560269c537a0
        Nexthop key opaque app data dump: TLV Type:32776, :bier-10-0.bier.0, Num of
entries:256, 1-3, 2-8091, 3-8103, 4-8092
        Kernel Table Id: 0
        State: <Active OpaqueData>
        Local AS: 200
        Age: 1:49:17
        Validation State: unverified
        Task: bier global task
        Announcement bits (1): 1-KRT
        AS path: I
        Statistics - Packets: 0, pps: 0, Bytes: 0
        Stats ID Group: Kernel ID = 4097, Stats IDs = { 4026531842 }
        Thread: junos-main

```

BFR-ID 2 and BFR-ID 3 are looked up in the BIFT :bier-10-0.bier.0 and their corresponding push labels, and next hops are determined. Issue the show route table :bier-10-0.bier.0 command to view this next hop information.

```

user@routerPE1> show route table :bier-10-0.bier.0

:bier-10-0.bier.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

1/16
    *[BIER/70] 01:56:21
        to table mpls.0
2/16
    *[BIER/70] 01:51:14
    > to 10.1.2.2 via et-0/0/11.0, Push 16
        to 10.1.2.2 via et-0/0/25.0, Push 16
3/16
    *[BIER/70] 01:51:06
    > to 10.1.2.2 via et-0/0/11.0, Push 16
        to 10.1.5.5 via et-0/0/9.0, Push 16
        to 10.1.2.2 via et-0/0/25.0, Push 16
4/16
    *[BIER/70] 01:50:36
    > to 10.1.5.5 via et-0/0/9.0, Push 16

```

To view the F-BM and its BIER neighbors, issue the `show route table :bier-10-0.bier.0 extensive` command.

```

user@routerPE1> show route table :bier-10-0.bier.0 extensive

:bier-10-0.bier.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
2/16 (1 entry, 0 announced)
    *BIER Preference: 70
        Next hop type: List, Next hop index: 8091
        Address: 0x560263986adc
        Next-hop reference count: 2, non-key opaque handle: 0x56026719b540
        Nexthop nonkey opaque app data dump: TLV type:32829, Flag:0x4
        Kernel Table Id: 0
        Next hop: ELNH Address 0x56026337c87c, selected
            Next hop type: Router, Next hop index: 8083
            Address: 0x56026337c87c
            Next-hop reference count: 2, key opaque handle: 0x560269c4eb80, non-key
opaque handle: 0x560269c50ce0
                Nexthop key opaque app data dump: TLV type:32794 APP data:Bier nbr
0x5602661cd418, label:16, Addr:10.2.2.2
                Nexthop nonkey opaque app data dump: TLV type:32829, Flag:0x4,
FBM:0:0:0:0:0:0:0:00000006
                    , Next-hop session id: 5
                    Kernel Table Id: 0
                    Next hop: 10.1.2.2 via et-0/0/11.0
                    Label operation: Push 16
                    Label TTL action: prop-ttl

```



```

Load balance label: Label 16: None;
Label element ptr: 0x560266949df0
Label parent element ptr: (nil)
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Statistics ID Group: Kernel ID = 4110, Stats IDs = { 4026531855 }
...trimmed

```

The F-BM can also be viewed by issuing the `show bier neighbor` command.

The output also displays Packets(pps). In the example output below, the packet counter for BIER neighbor 10.2.2.2 with next hop 10.1.2.2 shows 141753 while the other neighbors show null. This indicates the current forwarding path for multicast traffic. After masking the bit string with its F-BM, it is forwarded to its neighbors with push label 16.

```

user@routerPE1> show bier neighbor
BIER info:
Subdomain ID: 10   BSL: 256
BIER Neighbor          Label      Nexthop
Packets(pps)         Bytes
10.2.2.2              16         10.1.2.2
141753( 1)          18711396
FBM:0:0:0:0:0:0:0:0:00000006
10.2.2.2              16         10.1.2.2
0( 0)                 0
FBM:0:0:0:0:0:0:0:0:00000006
10.5.5.5              16         10.1.5.5
0( 0)                 0
FBM:0:0:0:0:0:0:0:0:0000000c

```

On PE2:

The incoming label is looked up in the local `mpls.0` table and the corresponding BIFT is determined. Issue the `show route table mpls.0 protocol bier label 16 extensive` command to view label 16's table information.

```

user@routerPE2> show route table mpls.0 protocol bier label 16 extensive

mpls.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
16 (1 entry, 1 announced)
TSI:
KRT in-kernel 16   /52 -> {[6021]}

```

```

Opaque data client: BIER
Opaque data: TLV type:32820 APP flag:0x80
Address: 0x560719d9aac0
Opaque-data reference count: 24
Stats ID Group: Kernel ID = 4097, Stats IDs = { 4026531842 }
    *BIER Preference: 70
        Next hop type: Multicast (IPv4) Composite, Next hop index: 6021
        Address: 0x56071658a084
        Next-hop reference count: 2, key opaque handle: 0x56071c81ffe0
        Nexthop key opaque app data dump: TLV Type:32776, :bier-10-0.bier.0, Num of
entries:256, 1-9093, 2-3, 3-9082, 4-9140
        Kernel Table Id: 0
        State: <Active OpaqueData>
        Local AS: 200
        Age: 1:55:50
        Validation State: unverified
        Task: bier global task
        Announcement bits (1): 1-KRT
        AS path: I
        Statistics - Packets: 0, pps: 0, Bytes: 0
        Stats ID Group: Kernel ID = 4097, Stats IDs = { 4026531842 }
        Thread: junos-main

```

Bits 2 and 3 are looked up in the BIFT :bier-10-0.bier.0. For bit 2, the next hop is the local mpls.0 table and for bit 3, the next hop is PE3 with push label 16. Issue the show route table :bier-10-0.bier.0 command to verify this information.

```

user@routerPE2> show route table :bier-10-0.bier.0

:bier-10-0.bier.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1/16
    *[BIER/70] 02:00:43
    > to 10.1.2.1 via et-0/0/11.0, Push 16
    to 10.1.2.1 via et-0/0/24.0, Push 16

2/16
    *[BIER/70] 02:05:35
    to table mpls.0

3/16
    *[BIER/70] 02:00:48
    > to 10.2.3.3 via et-0/0/12.0, Push 16

```

4/16

```
*[BIER/70] 02:00:04
> to 10.2.3.3 via et-0/0/12.0, Push 16
  to 10.2.5.5 via et-0/0/9.0, Push 16
```

For bit 2, the vrf label 990001 is looked up in the mpls.0 table to find its corresponding vrf.

```
root@user@routerPE2> show route table mpls.0 label 990001

mpls.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

990001          *[VPN/0] 02:08:30
                > via lsi.0 (vrf1), Pop
```

Issue the show multicast route instance vrf1 extensive command to view the route.

```
user@routerPE2> show multicast route instance vrf1 extensive
Instance: vrf1 Family: INET

Group: 232.252.1.1
  Source: 172.11.21.21/32
  Upstream interface: lsi.0
  Downstream interface list:
    et-0/0/10.1
  Number of outgoing interfaces: 1
  Session description: Unknown
  Statistics: 0 kBps, 0 pps, 0 packets
  Next-hop ID: 6017
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:18:11
  Sensor ID: 0xf0000017
```

ON PE3:

Similar to PE2, the incoming label is looked up in the local mpls.0 table and the corresponding BIFT is determined. Issue the `show route table mpls.0 protocol bier label 16 extensive` command to view label 16's table information.

```

user@routerPE3> show route table mpls.0 label 16 extensive

mpls.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
16 (1 entry, 1 announced)
TSI:
KRT in-kernel 16 /52 -> {[6013]}
Opaque data client: BIER
Opaque data: TLV type:32820 APP flag:0x80
Address: 0x55c74399b1e0
Opaque-data reference count: 24
Stats ID Group: Kernel ID = 4097, Stats IDs = { 4026531842 }
  *BIER Preference: 70
    Next hop type: Multicast (IPv4) Composite, Next hop index: 6013
    Address: 0x55c740189184
    Next-hop reference count: 2, key opaque handle: 0x55c74642ea20
    Nexthop key opaque app data dump: TLV Type:32776, :bier-10-0.bier.0, Num of
entries:256, 1-9088, 2-9085, 3-3, 4-9116
    Kernel Table Id: 0
    State: <Active OpaqueData>
    Local AS: 200
    Age: 2:04:31
    Validation State: unverified
    Task: bier global task
    Announcement bits (1): 1-KRT
    AS path: I
    Statistics - Packets: 0, pps: 0, Bytes: 0
    Stats ID Group: Kernel ID = 4097, Stats IDs = { 4026531842 }
    Thread: junos-main

```

Issue the `show route table :bier-10-0.bier.0` command to view next hop information. Bit 3 is looked up in the BIFT and the next hop is the local mpls.0 table.

```

user@routerPE3> show route table :bier-10-0.bier.0

:bier-10-0.bier.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

1/16
    *[BIER/70] 02:06:09
    > to 10.3.5.5 via ae0.0, Push 16
      to 10.2.3.2 via et-0/0/11.0, Push 16
2/16
    *[BIER/70] 02:06:15
    > to 10.2.3.2 via et-0/0/11.0, Push 16
3/16
    *[BIER/70] 02:10:46
    to table mpls.0
4/16
    *[BIER/70] 02:05:52
    > to 10.3.4.4 via et-0/0/13.0, Push 16

```

The vrf label 990001 is looked up in the mpls.0 table, to find the vrf for bit 3.

```

user@routerPE3> show route table mpls.0 label 990001

mpls.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

990001          *[VPN/0] 02:11:39
                > via lsi.0 (vrf1), Pop

```

Issue the show multicast route instance vrf1 extensive command to view the route in the vrf.

```

user@routerPE3> show multicast route instance vrf1 extensive
Instance: vrf1 Family: INET

Group: 232.252.1.1
  Source: 172.11.21.21/32
  Upstream interface: lsi.0
  Downstream interface list:
    et-0/0/10.1
  Number of outgoing interfaces: 1
  Session description: Unknown
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 6017
  Upstream protocol: MVPN
  Route state: Active
  Forwarding state: Forwarding

```

```
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
Uptime: 00:21:33
Sensor ID: 0xf0000015
```

RELATED DOCUMENTATION

bier (Protocols)

bier-sub-domain (Protocols IS-IS)

show route table

Prevent Routing Loops with Reverse Path Forwarding

IN THIS CHAPTER

- [Examples: Configuring Reverse Path Forwarding | 1149](#)

Examples: Configuring Reverse Path Forwarding

IN THIS SECTION

- [Understanding Multicast Reverse Path Forwarding | 1149](#)
- [Multicast RPF Configuration Guidelines | 1151](#)
- [Example: Configuring a Dedicated PIM RPF Routing Table | 1152](#)
- [Example: Configuring a PIM RPF Routing Table | 1157](#)
- [Example: Configuring RPF Policies | 1164](#)
- [Example: Configuring PIM RPF Selection | 1167](#)

Understanding Multicast Reverse Path Forwarding

IN THIS SECTION

- [RPF Table | 1151](#)

RPF Table

The RPF table plays the key role in the multicast router. The RPF table is consulted for every RPF check, which is performed at intervals on multicast packets entering the multicast router. Distribution trees of all types rely on the RPF table to form properly, and the multicast forwarding state also depends on the RPF table.

RPF checks are performed only on unicast addresses to find the upstream interface for the multicast source or RP.

The routing table used for RPF checks can be the same routing table used to forward unicast IP packets, or it can be a separate routing table used only for multicast RPF checks. In either case, the RPF table contains only unicast routes, because the RPF check is performed on the source address of the multicast packet, not the multicast group destination address, and a multicast address is forbidden from appearing in the source address field of an IP packet header. The unicast address can be used for RPF checks because there is only one source host for a particular stream of IP multicast content for a multicast group address, although the same content could be available from multiple sources.

If the same routing table used to forward unicast packets is also used for the RPF checks, the routing table is populated and maintained by the traditional unicast routing protocols such as BGP, IS-IS, OSPF, and the Routing Information Protocol (RIP). If a dedicated multicast RPF table is used, this table must be populated by some other method. Some multicast routing protocols (such as the Distance Vector Multicast Routing Protocol [DVMRP]) essentially duplicate the operation of a unicast routing protocol and populate a dedicated RPF table. Others, such as PIM, do not duplicate routing protocol functions and must rely on some other routing protocol to set up this table, which is why PIM is protocol independent. .

Some traditional routing protocols such as BGP and IS-IS now have extensions to differentiate between different sets of routing information sent between routers for unicast and multicast. For example, there is multiprotocol BGP (MBGP) and multiprotocol routing in IS-IS (M-IS-IS). IS-IS routes can be added to the RPF table even when special features such as traffic engineering and “shortcuts” are turned on. Multicast Open Shortest Path First (MOSPF) also extends OSPF for multicast use, but goes further than MBGP or M-IS-IS and makes MOSPF into a complete multicast routing protocol on its own. When these routing protocols are used, routes can be tagged as multicast RPF routers and used by the receiving router differently than the unicast routing information.

Using the main unicast routing table for RPF checks provides simplicity. A dedicated routing table for RPF checks allows a network administrator to set up separate paths and routing policies for unicast and multicast traffic, allowing the multicast network to function more independently of the unicast network.

Multicast RPF Configuration Guidelines

You use multicast RPF checks to prevent multicast routing loops. Routing loops are particularly debilitating in multicast applications because packets are replicated with each pass around the routing loop.

In general, a router is to forward a multicast packet only if it arrives on the interface closest (as defined by a unicast routing protocol) to the origin of the packet, whether source host or rendezvous point (RP). In other words, if a unicast packet would be sent to the “destination” (the reverse path) on the interface that the multicast packet arrived on, the packet passes the RPF check and is processed. Multicast (or unicast) packets that fail the RPF check are not forwarded (this is the default behavior). For an overview of how a Juniper Networks router implements RPF checks with tables, see [Understanding Multicast Reverse Path Forwarding](#).

However, there are network router configurations where multicast packets that fail the RPF check need to be forwarded. For example, when point-to-multipoint label-switched paths (LSPs) are used for distributing multicast traffic to PIM “islands” downstream from the egress router, the interface on which the multicast traffic arrives is not always the RPF interface. This is because LSPs do not follow the normal next-hop rules of independent packet routing.

In cases such as these, you can configure policies on the PE router to decide which multicast groups and sources are exempt from the default RPF check.

SEE ALSO

[Junos OS MPLS Applications User Guide](#)

[Routing Policies, Firewall Filters, and Traffic Policers User Guide](#)

Example: Configuring a Dedicated PIM RPF Routing Table

IN THIS SECTION

- [Requirements | 1152](#)
- [Overview | 1153](#)
- [Configuration | 1154](#)

This example explains how to configure a dedicated Protocol Independent Multicast (PIM) reverse path forwarding (RPF) routing table.

Requirements

Before you begin:

- Configure the router interfaces. See the [Interfaces User Guide for Security Devices](#).

- Enable PIM. See ["PIM Overview" on page 275](#).

This example uses the following software components:

- Junos OS Release 7.4 or later

Overview

By default, PIM uses the **inet.0** routing table as its RPF routing table. PIM uses an RPF routing table to resolve its RPF neighbor for a particular multicast source address and to resolve the RPF neighbor for the rendezvous point (RP) address. PIM can optionally use **inet.2** as its RPF routing table. The **inet.2** routing table is dedicated to this purpose.

PIM uses a single routing table for its RPF check, this ensures that the route with the longest matching prefix is chosen as the RPF route.

If multicast routes are exchanged by Multiprotocol Border Gateway Protocol MP-BGP or multitopology IS-IS, they are placed in **inet.2** by default.

Using **inet.2** as the RPF routing table enables you to have a control plane for multicast, which is independent of the normal unicast routing table. You might want to use **inet.2** as the RPF routing table for any of the following reasons:

- If you use traffic engineering or have an interior gateway protocol (IGP) configured for shortcuts, the router has label-switched paths (LSPs) installed as the next hops in **inet.2**. By applying policy, you can have the router install the routes with non-MPLS next-hops in the **inet.2** routing table.
- If you have an MPLS network that does not support multicast traffic over LSP tunnels, you need to configure the router to use a routing table other than **inet.0**. You can have the **inet.2** routing table populated with native IGP, BGP, and interface routes that can be used for RPF.

To populate the PIM RPF table, you use rib groups. A rib group is defined with the `rib-groups` statement at the `[edit routing-options]` hierarchy level. The rib group is applied to the PIM protocol by including the `rib-group` statement at the `[edit pim]` hierarchy level. A rib group is most frequently used to place routes in multiple routing tables.

When you configure rib groups for PIM, keep the following in mind:

- The `import-rib` statement copies routes from the protocol to the routing table.
- The `export-rib` statement has no effect on PIM. However the same rib-group can be applied to BGP, which uses the `export-rib` table to source routes to advertise to peers in the case of inter-domain networking.
- Only the first rib routing table specified in the `import-rib` statement is used by PIM for RPF checks.

You can also configure IS-IS or OSPF to populate **inet.2** with routes that have regular IP next hops. This allows RPF to work properly even when MPLS is configured for traffic engineering, or when IS-IS or OSPF are configured to use “shortcuts” for local traffic.

You can also configure the PIM protocol to use a rib group for RPF checks under a virtual private network (VPN) routing instance. In this case the rib group is still defined at the [edit routing-options] hierarchy level.

Configuration

IN THIS SECTION

- [Configuring a PIM RPF Routing Table Group Using Interface Routes | 1154](#)
- [Verifying Multicast RPF Table | 1156](#)

Configuring a PIM RPF Routing Table Group Using Interface Routes

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.

```
set routing-options rib-groups mcast-rpf-rib import-rib inet.2
set protocols pim rib-group mcast-rpf-rib
set routing-options interface-routes rib-group inet if-rib
set routing-options rib-groups if-rib import-rib [ inet.0 inet.2 ]
```

Step-by-Step Procedure

In this example, the network administrator has decided to use the **inet.2** routing table for RPF checks. In this process, local routes are copied into this table by using an interface rib group.

To define an interface routing table group and use it to populate **inet.2** for RPF checks:

1. Use the `show multicast rpf` command to verify that the multicast RPF table is not populated with routes.

```
user@host> show multicast rpf
instance is not running
```

2. Create a multicast routing table group named **mcast-rpf-rib**.

Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the `import-rib` statement).

Include the `import-rib` statement and specify the **inet.2** routing table at the `[edit routing-options rib-groups]` hierarchy level.

```
[edit routing-options rib-groups]
user@host# set mcast-rpf-rib import-rib inet.2
```

3. Configure PIM to use the **mcast-rpf-rib** rib group.

The rib group for PIM can be applied globally or in a routing instance. In this example, the global configuration is shown.

Include the `rib-group` statement and specify the **mcast-rpf-rib** rib group at the `[edit protocols pim]` hierarchy level.

```
[edit protocols pim]
user@host# set rib-group mcast-rpf-rib
```

4. Create an interface rib group named **if-rib**.

Include the `rib-group` statement and specify the **inet** address family at the `[edit routing-options interface-routes]` hierarchy level.

```
[edit routing-options interface-routes]
user@host# set rib-group inet if-rib
```

5. Configure the **if-rib** rib group to import routes from the **inet.0** and **inet.2** routing tables.

Include the `import-rib` statement and specify the **inet.0** and **inet.2** routing tables at the `[edit routing-options rib-groups]` hierarchy level.

```
[edit routing-options rib-groups]
user@host# set if-rib import-rib [ inet.0 inet.2 ]
```

6. Commit the configuration.

```
user@host# commit
```

Verifying Multicast RPF Table

Purpose

Verify that the multicast RPF table is now populated with routes.

Action

Use the `show multicast rpf` command.

```
user@host> show multicast rpf
Multicast RPF table: inet.2 , 10 entries

10.0.24.12/30
  Protocol: Direct
  Interface: fe-0/1/2.0

10.0.24.13/32
  Protocol: Local

10.0.27.12/30
  Protocol: Direct
  Interface: fe-0/1/3.0

10.0.27.13/32
  Protocol: Local

10.0.224.8/30
  Protocol: Direct
```

```
Interface: ge-1/3/3.0

10.0.224.9/32
  Protocol: Local

127.0.0.1/32
  Inactive

192.168.2.1/32
  Protocol: Direct
  Interface: lo0.0

192.168.187.0/25
  Protocol: Direct
  Interface: fxp0.0

192.168.187.12/32
  Protocol: Local
```

Meaning

The first line of the sample output shows that the **inet.2** table is being used and that there are 10 routes in the table. The remainder of the sample output lists the routes that populate the **inet.2** routing table.

SEE ALSO

Example: Enabling OSPF Traffic Engineering Support

[traffic-engineering](#)

show multicast rpf

Example: Configuring a PIM RPF Routing Table

IN THIS SECTION

- [Requirements | 1158](#)
- [Overview | 1158](#)
- [Configuration | 1159](#)
- [Verification | 1161](#)

This example shows how to configure and apply a PIM RPF routing table.

Requirements

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See ["Configuring the Session Announcement Protocol" on page 573](#).
8. Configure IGMP. See ["Configuring IGMP" on page 25](#).
9. Configure the PIM static RP. See ["Configuring Static RP" on page 342](#).
10. Filter PIM register messages from unauthorized groups and sources. See ["Example: Rejecting Incoming PIM Register Messages on RP Routers" on page 389](#) and ["Example: Stopping Outgoing PIM Register Messages on a Designated Router" on page 382](#).

Overview

In this example, you name the new RPF routing table group **multicast-rfp-rib** and use **inet.2** for its export as well as its import routing table. Then you create a routing table group for the interface routes and name the RPF **if-rib**. Finally, you use **inet.2** and **inet.0** for its import routing tables, and add the new interface routing table group to the interface routes.

Configuration

IN THIS SECTION

- [Procedure | 1159](#)

Procedure

CLI Quick Configuration

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

```
set routing-options rib-groups multicast-rpf-rib export-rib inet.2
set routing-options rib-groups multicast-rpf-rib import-rib inet.2
set protocols pim rib-group multicast-rpf-rib
set routing-options rib-groups if-rib import-rib inet.2
set routing-options rib-groups if-rib import-rib inet.0
set routing-options interface-routes rib-group inet if-rib
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [Using the CLI Editor in Configuration Mode](#) in the [Junos OS CLI User Guide](#).

To configure the PIM RPF routing table:

1. Configure a routing option and a group.

```
[edit]
user@host# edit routing-options rib-groups
```

2. Configure a name.

```
[edit routing-options rib-groups]
user@host# set multicast-rpf-rib export-rib inet.2
```

3. Create a new group for the RPF routing table.

```
[edit routing-options rib-groups]
user@host# set multicast-rpf-rib import-rib inet.2
```

4. Apply the new RPF routing table.

```
[edit protocols pim]
user@host# set rib-group multicast-rpf-rib
```

5. Create a routing table group for the interface routes.

```
[edit]
user@host# edit routing-options rib-groups
```

6. Configure a name for import routing table.

```
[edit routing-options rib-groups]
user@host# set if-rib import-rib inet.2
user@host# set if-rib import-rib inet.0
```

7. Set group to interface routes.

```
[edit routing-options interface-routes]
user@host# set rib-group inet if-rib
```

Results

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

```
[edit]
user@host# show protocols
pim {
  rib-group inet multicast-rpf-rib;
}
[edit]
user@host# show routing-options
interface-routes {
  rib-group inet if-rib;
}
static {
  route 0.0.0.0/0 next-hop 10.100.37.1;
}
rib-groups {
  multicast-rpf-rib {
    export-rib inet.2;
    import-rib inet.2;
  }
  if-rib {
    import-rib [ inet.2 inet.0 ];
  }
}
```

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

Verification

IN THIS SECTION

- [Verifying SAP and SDP Addresses and Ports | 1162](#)
- [Verifying the IGMP Version | 1162](#)
- [Verifying the PIM Mode and Interface Configuration | 1163](#)
- [Verifying the PIM RP Configuration | 1163](#)

- [Verifying the RPF Routing Table Configuration | 1163](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying SAP and SDP Addresses and Ports

Purpose

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

Action

From operational mode, enter the `show sap listen` command.

Verifying the IGMP Version

Purpose

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

Action

From operational mode, enter the `show igmp interface` command.

```
user@host> show igmp interface
Interface: ge-0/0/0.0
  Querier: 192.168.4.36
  State:      Up Timeout: 197 Version: 2 Groups: 0

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

Derived Parameters:
```

```
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

Verifying the PIM Mode and Interface Configuration

Purpose

Verify that PIM sparse mode is configured on all applicable interfaces.

Action

From operational mode, enter the `show pim interfaces` command.

Verifying the PIM RP Configuration

Purpose

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

Action

From operational mode, enter the `show pim rps` command.

Verifying the RPF Routing Table Configuration

Purpose

Verify that the PIM RPF routing table is configured correctly.

Action

From operational mode, enter the `show multicast rpf` command.

SEE ALSO

[Configuring PIM Filtering | 376](#)

[Example: Configuring a Dedicated PIM RPF Routing Table | 1152](#)

[Multicast Configuration Overview | 17](#)

[Verifying a Multicast Configuration](#)

Example: Configuring RPF Policies

IN THIS SECTION

- Requirements | 1164
- Overview | 1164
- Configuration | 1165
- Verification | 1167

A multicast RPF policy disables RPF checks for a particular multicast (S,G) pair. You usually disable RPF checks on egress routing devices of a point-to-multipoint label-switched path (LSP), because the interface receiving the multicast traffic on a point-to-multipoint LSP egress router might not always be the RPF interface.

This example shows how to configure an RPF check policy named **disable-RPF-on-PE**. The **disable-RPF-on-PE** policy disables RPF checks on packets arriving for group 228.0.0.0/8 or from source address 196.168.25.6.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview

An RPF policy behaves like an import policy. If no policy term matches the input packet, the default action is to accept (that is, to perform the RPF check). The **route-filter** statement filters group addresses, and the **source-address-filter** statement filters source addresses.

This example shows how to configure each condition as a separate policy and references both policies in the **rpf-check-policy** statement. This allows you to associate groups in one policy and sources in the other.



NOTE: Be careful when disabling RPF checks on multicast traffic. If you disable RPF checks in some configurations, multicast loops can result.

Changes to an RPF check policy take effect immediately:

- If no policy was previously configured, the policy takes effect immediately.
- If the policy name is changed, the new policy takes effect immediately and any packets no longer filtered are subjected to the RPF check.
- If the policy is deleted, all packets formerly filtered are subjected to the RPF check.
- If the underlying policy is changed, but retains the same name, the new conditions take effect immediately and any packets no longer filtered are subjected to the RPF check.

Configuration

IN THIS SECTION

- [Procedure | 1165](#)
- [Results | 1166](#)

Procedure

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.

```
set policy-options policy-statement disable-RPF-from-group term first from route-filter
228.0.0.0/8 orlonger
set policy-options policy-statement disable-RPF-from-group term first then reject
set policy-options policy-statement disable-RPF-from-source term first from source-address-
filter 192.168.25.6/32 exact
set policy-options policy-statement disable-RPF-from-source term first then reject
set routing-options multicast rpf-check-policy [ disable-RPF-from-group disable-RPF-from-source ]
```

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 [Junos OS CLI User Guide](#).

To configure an RPF policy:

1. Configure a policy for group addresses.

```
[edit policy-options]
user@host# set policy-statement disable-RPF-for-group term first from route-filter
228.0.0.0/8 orlonger
user@host# set policy-statement disable-RPF-for-group term first then reject
```

2. Configure a policy for a source address.

```
[edit policy-options]
user@host# set policy-statement disable-RPF-for-source term first from source-address-filter
192.168.25.6/32 exact
user@host# set policy-statement disable-RPF-for-source term first then reject
```

3. Apply the policies.

```
[edit routing-options]
user@host# set multicast rpf-check-policy [ disable-RPF-for-group disable-RPF-for-source ]
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** and **show routing-options** commands.

```
user@host# show policy-options
policy-statement disable-RPF-from-group {
  term first {
```



```
        from {
            route-filter 228.0.0.0/8 orlonger;
        }
        then reject;
    }
}
policy-statement disable-RPF-from-source {
    term first {
        from {
            source-address-filter 192.168.25.6/32 exact;
        }
        then reject;
    }
}
```

```
user@host# show routing-options
multicast {
    rpf-check-policy [ disable-RPF-from-group disable-RPF-from-source ];
}
```

Verification

To verify the configuration, run the **show multicast rpf** command.

SEE ALSO

[Example: Configuring Ingress PE Redundancy | 1327](#)

Example: Configuring PIM RPF Selection

IN THIS SECTION

- [Requirements | 1168](#)
- [Overview | 1168](#)
- [Configuration | 1170](#)
- [Verification | 1172](#)

This example shows how to configure and verify the multicast PIM RPF next-hop neighbor selection for a group or (S,G) pair.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Make sure that the RPF next-hop neighbor you want to specify is operating.

Overview

IN THIS SECTION

- [Topology | 1169](#)

Multicast PIM RPF neighbor selection allows you to specify the RPF neighbor (next hop) and source address for a single group or multiple groups using a prefix list. RPF neighbor selection can only be configured for VPN routing and forwarding (VRF) instances.

If you have multiple service VRFs through which a receiver VRF can learn the same source or rendezvous point (RP) address, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows. However, if RPF neighbor selection is configured, RPF checks are based on your configuration instead of the unicast routing protocols.

You can use this static RPF selection as a building block for particular applications. For example, an extranet. Suppose you want to split the multicast flows among parallel PIM links or assign one multicast flow to a specific PIM link. With static RPF selection configured, the router sends join and prune messages based on the configuration.

You can use wildcards to designate the source address. Whether or not you use wildcards affects how the PIM joins work:

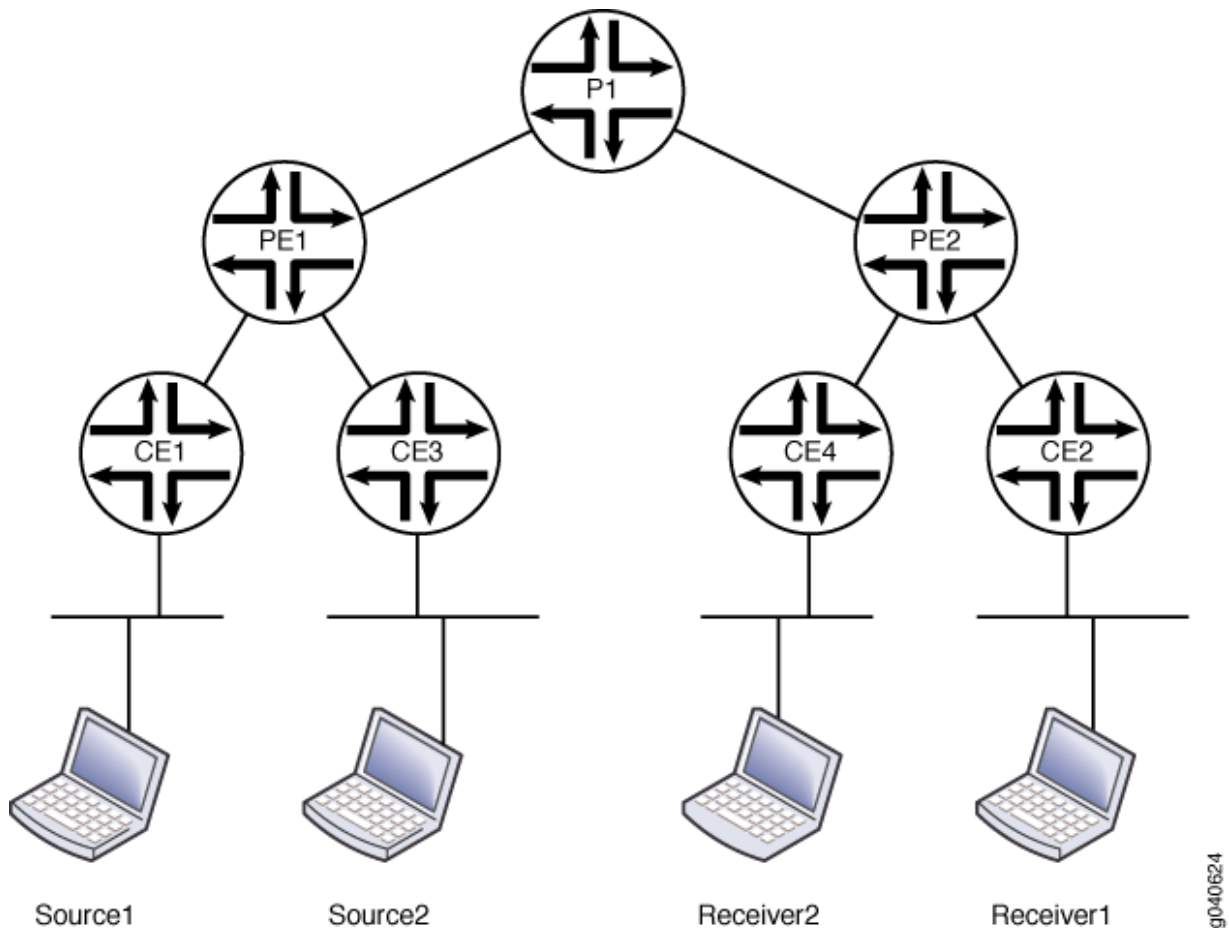
- If you configure only a source prefix for a group, all (*,G) joins are sent to the next-hop neighbor selected by the unicast protocol, while (S,G) joins are sent to the next-hop neighbor specified for the source.

- If you configure only a wildcard source for a group, all (*,G) and (S,G) joins are sent to the upstream interface pointing to the wildcard source next-hop neighbor.
- If you configure both a source prefix and a wildcard source for a group, all (S,G) joins are sent to the next-hop neighbor defined for the source prefix, while (*,G) joins are sent to the next-hop neighbor specified for the wildcard source.

Topology

Figure 135 on page 1169 shows the topology used in this example.

Figure 135: PIM RPF Selection



In this example, the RPF selection is configured on the receiver provider edge router (PE2).

Configuration

IN THIS SECTION

- [Procedure | 1170](#)

Procedure

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.

```
set routing-instance vpn-a protocols pim rpf-selection group 225.5.0.0/16 wildcard-source next-hop 10.12.5.2
set routing-instance vpn-a protocols pim rpf-selection prefix-list group12 wildcard-source next-hop 10.12.31.2
set routing-instance vpn-a protocols pim rpf-selection prefix-list group34 source 22.1.12.0/24 next-hop 10.12.32.2
set policy-options prefix-list group12 225.1.1.0/24
set policy-options prefix-list group12 225.2.0.0/16
set policy-options prefix-list group34 225.3.3.3/32
set policy-options prefix-list group34 225.4.4.0/24
```

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 [Junos OS CLI User Guide](#).

To configure PIM RPF selection:

1. On PE2, configure RFP selection in a routing instance.

```
[edit routing-instance vpn-a protocols pim]
user@host# set rpf-selection group 225.5.0.0/16 wildcard-source next-hop 10.12.5.2
user@host# set rpf-selection prefix-list group12 wildcard-source next-hop 10.12.31.2
```

```
user@host# set rpf-selection prefix-list group34 source 22.1.12.0/24 next-hop 10.12.32.2
user@host# exit
```

2. On PE2, configure the policy.

```
[edit policy-options]
set prefix-list group12 225.1.1.0/24
set prefix-list group12 225.2.0.0/16
set prefix-list group34 225.3.3.3/32
set prefix-list group34 225.4.4.0/24
```

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

```
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the **show policy-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@host# show policy-options
prefix-list group12 {
    225.1.1.0/24;
    225.2.0.0/16;
}
prefix-list group34 {
    225.3.3.3/32;
    225.4.4.0/24;
}
```

```
user@host# show routing-instances
vpn-a{
    protocols {
        pim {
            rpf-selection {
                group 225.5.0.0/16 {
                    wildcard-source {
```

```
        next-hop 10.12.5.2;
    }
}
prefix-list group12 {
    wildcard-source {
        next-hop 10.12.31.2;
    }
}
prefix-list group34 {
    source 22.1.12.0/24 {
        next-hop 10.12.32.2;
    }
}
}
}
}
}
```

Verification

To verify the configuration, run the following commands, checking the upstream interface and the upstream neighbor:

- **show pim join extensive**
- **show multicast route**

RELATED DOCUMENTATION

| [Example: Configuring Ingress PE Redundancy](#) | 1327

Use Multicast-Only Fast Reroute (MoFRR) to Minimize Packet Loss During Link Failures

IN THIS CHAPTER

- [Understanding Multicast-Only Fast Reroute | 1173](#)
- [Configuring Multicast-Only Fast Reroute | 1182](#)
- [Example: Configuring Multicast-Only Fast Reroute in a PIM Domain | 1185](#)
- [Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches | 1197](#)
- [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1208](#)

Understanding Multicast-Only Fast Reroute

IN THIS SECTION

- [MoFRR Overview | 1174](#)
- [PIM Functionality | 1176](#)
- [Multipoint LDP Functionality | 1177](#)
- [Packet Forwarding | 1178](#)
- [Limitations and Caveats | 1180](#)

Multicast-only fast reroute (MoFRR) minimizes packet loss for traffic in a multicast distribution tree when link failures occur, enhancing multicast routing protocols like Protocol Independent Multicast (PIM) and multipoint Label Distribution Protocol (multipoint LDP) on devices that support these features.



NOTE: On switches, MoFRR with MPLS label-switched paths and multipoint LDP is not supported.

For MX Series routers, MoFRR is supported only on MX Series routers with MPC line cards. As a prerequisite, you must configure the router into `network-services enhanced-ip` mode, and all the line cards in the router must be MPCs.

With MoFRR enabled, devices send join messages on primary and backup upstream paths toward a multicast source. Devices receive data packets from both the primary and backup paths, and discard the redundant packets based on priority (weights that are assigned to the primary and backup paths). When a device detects a failure on the primary path, it immediately starts accepting packets from the secondary interface (the backup path). The fast switchover greatly improves convergence times upon primary path link failures.

One application for MoFRR is streaming IPTV. IPTV streams are multicast as UDP streams, so any lost packets are not retransmitted, leading to a less-than-satisfactory user experience. MoFRR can improve the situation.

MoFRR Overview

With fast reroute on unicast streams, an upstream routing device preestablishes MPLS label-switched paths (LSPs) or precomputes an IP loop-free alternate (LFA) fast reroute backup path to handle failure of a segment in the downstream path.

In multicast routing, the receiving side usually originates the traffic distribution graphs. This is unlike unicast routing, which generally establishes the path from the source to the receiver. PIM (for IP), multipoint LDP (for MPLS), and RSVP-TE (for MPLS) are protocols that are capable of establishing multicast distribution graphs. Of these, PIM and multipoint LDP receivers initiate the distribution graph setup, so MoFRR can work with these two multicast protocols where they are supported.

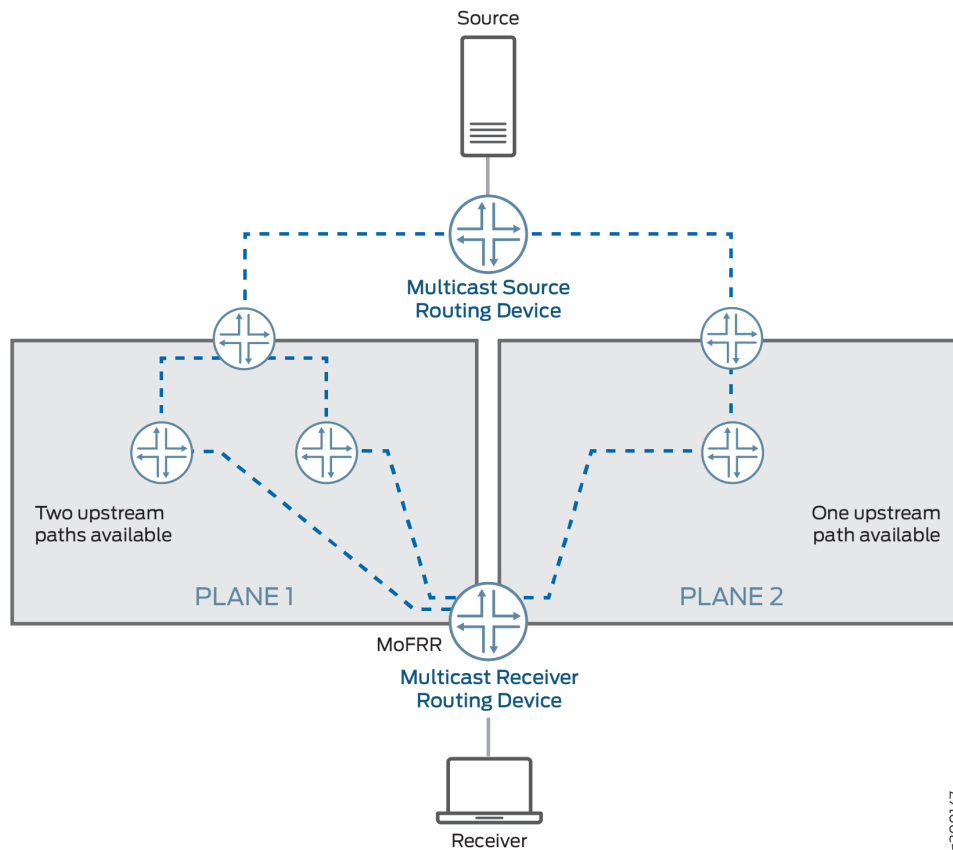
In a multicast tree, if the device detects a network component failure, it takes some time to perform a reactive repair, leading to significant traffic loss while setting up an alternate path. MoFRR reduces traffic loss in a multicast distribution tree when a network component fails. With MoFRR, one of the downstream routing devices sets up an alternative path toward the source to receive a backup live stream of the same multicast traffic. When a failure happens along the primary stream, the MoFRR routing device can quickly switch to the backup stream.

With MoFRR enabled, for each (S,G) entry, the device uses two of the available upstream interfaces to send a join message and to receive multicast traffic. The protocol attempts to select two disjoint paths if two such paths are available. If disjoint paths are not available, the protocol selects two non-disjoint paths. If two non-disjoint paths are not available, only a primary path is selected with no backup. MoFRR prioritizes the disjoint backup in favor of load balancing the available paths.

MoFRR is supported for both IPv4 and IPv6 protocol families.

Figure 136 on page 1175 shows two paths from the multicast receiver routing device (also referred to as the egress provider edge (PE) device) to the multicast source routing device (also referred to as the ingress PE device).

Figure 136: MoFRR Sample Topology



With MoFRR enabled, the egress (receiver side) routing device sets up two multicast trees, a primary path and a backup path, toward the multicast source for each (S,G). In other words, the egress routing device propagates the same (S,G) join messages toward two different upstream neighbors, thus creating two multicast trees.

One of the multicast trees goes through plane 1 and the other through plane 2, as shown in Figure 136 on page 1175. For each (S,G), the egress routing device forwards traffic received on the primary path and drops traffic received on the backup path.

MoFRR is supported on both equal-cost multipath (ECMP) paths and non-ECMP paths. The device needs to enable unicast loop-free alternate (LFA) routes to support MoFRR on non-ECMP paths. You enable LFA routes using the `link-protection` statement in the interior gateway protocol (IGP)

configuration. When you enable link protection on an OSPF or IS-IS interface, the device creates a backup LFA path to the primary next hop for all destination routes that traverse the protected interface.

Junos OS implements MoFRR in the IP network for IP MoFRR and at the MPLS label-edge routing device (LER) for multipoint LDP MoFRR.

Multipoint LDP MoFRR is used at the egress device of an MPLS network, where the packets are forwarded to an IP network. With multipoint LDP MoFRR, the device establishes two paths toward the upstream PE routing device for receiving two streams of MPLS packets at the LER. The device accepts one of the streams (the primary), and the other one (the backup) is dropped at the LER. If the primary path fails, the device accepts the backup stream instead. Inband signaling support is a prerequisite for MoFRR with multipoint LDP (see [Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs](#)).

PIM Functionality

Junos OS supports MoFRR for shortest-path tree (SPT) joins in PIM source-specific multicast (SSM) and any-source multicast (ASM). MoFRR is supported for both SSM and ASM ranges. To enable MoFRR for (*,G) joins, include the `mofrr-asm-starg` configuration statement at the [edit routing-options multicast stream-protection] hierarchy. For each group G, MoFRR will operate for either (S,G) or (*,G), but not both. (S,G) always takes precedence over (*,G).

With MoFRR enabled, a PIM routing device propagates join messages on two upstream reverse-path forwarding (RPF) interfaces to receive multicast traffic on both links for the same join request. MoFRR gives preference to two paths that do not converge to the same immediate upstream routing device. PIM installs appropriate multicast routes with upstream RPF next hops with two interfaces (for the primary and backup paths).

When the primary path fails, the backup path is upgraded to primary status, and the device forwards traffic accordingly. If there are alternate paths available, MoFRR calculates a new backup path and updates or installs the appropriate multicast route.

You can enable MoFRR with PIM join load balancing (see the `join-load-balance` automatic statement). However, in that case the distribution of join messages among the links might not be even. When a new ECMP link is added, join messages on the primary path are redistributed and load-balanced. The join messages on the backup path might still follow the same path and might not be evenly redistributed.

You enable MoFRR using the `stream-protection` configuration statement at the [edit routing-options multicast] hierarchy. MoFRR is managed by a set of filter policies.

When an egress PIM routing device receives a join message or an IGMP report, it checks for an MoFRR configuration and proceeds as follows:

- If the MoFRR configuration is not present, PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 136 on page 1175](#)).

- If the MoFRR configuration is present, the device checks for a policy configuration.
- If a policy is not present, the device checks for primary and backup paths (upstream interfaces), and proceeds as follows:
 - If primary and backup paths are not available—PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 136 on page 1175](#)).
 - If primary and backup paths are available—PIM sends the join message upstream toward two of the available upstream neighbors. Junos OS sets up primary and secondary multicast paths to receive multicast traffic (for example, plane 1 in [Figure 136 on page 1175](#)).
- If a policy is present, the device checks whether the policy allows MoFRR for this (S,G), and proceeds as follows:
 - If this policy check fails—PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 136 on page 1175](#)).
 - If this policy check passes—The device checks for primary and backup paths (upstream interfaces).
 - If the primary and backup paths are not available, PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 136 on page 1175](#)).
 - If the primary and backup paths are available, PIM sends the join message upstream toward two of the available upstream neighbors. The device sets up primary and secondary multicast paths to receive multicast traffic (for example, plane 1 in [Figure 136 on page 1175](#)).

Multipoint LDP Functionality

To avoid MPLS traffic duplication, multipoint LDP usually selects only one upstream path. (See section 2.4.1.1. Determining One's 'upstream LSR' in RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*.)

For multipoint LDP with MoFRR, the multipoint LDP device selects two separate upstream peers and sends two separate labels, one to each upstream peer. The device uses the same algorithm described in RFC 6388 to select the primary upstream path. The device uses the same algorithm to select the backup upstream path but excludes the primary upstream LSR as a candidate. The two different upstream peers send two streams of MPLS traffic to the egress routing device. The device selects only one of the upstream neighbor paths as the primary path from which to accept the MPLS traffic. The other path becomes the backup path, and the device drops that traffic. When the primary upstream path fails, the device starts accepting traffic from the backup path. The multipoint LDP device selects the two upstream paths based on the interior gateway protocol (IGP) root device next hop.

A *forwarding equivalency class (FEC)* is a group of IP packets that are forwarded in the same manner, over the same path, and with the same forwarding treatment. Normally, the label that is put on a particular packet represents the FEC to which that packet is assigned. In MoFRR, two routes are placed

into the `mpls.0` table for each FEC—one route for the primary label and the other route for the backup label.

If there are parallel links toward the same immediate upstream device, the device considers both parallel links to be the primary. At any point in time, the upstream device sends traffic on only one of the multiple parallel links.

A *bud node* is an LSR that is an egress LSR, but also has one or more directly connected downstream LSRs. For a bud node, the traffic from the primary upstream path is forwarded to a downstream LSR. If the primary upstream path fails, the MPLS traffic from the backup upstream path is forwarded to the downstream LSR. This means that the downstream LSR next hop is added to both MPLS routes along with the egress next hop.

As with PIM, you enable MoFRR with multipoint LDP using the `stream-protection` configuration statement at the `[edit routing-options multicast]` hierarchy, and it's managed by a set of filter policies.

If you have enabled the multipoint LDP point-to-multipoint FEC for MoFRR, the device factors the following considerations into selecting the upstream path:

- The targeted LDP sessions are skipped if there is a nontargeted LDP session. If there is a single targeted LDP session, the targeted LDP session is selected, but the corresponding point-to-multipoint FEC loses the MoFRR capability because there is no interface associated with the targeted LDP session.
- All interfaces that belong to the same upstream LSR are considered to be the primary path.
- For any root-node route updates, the upstream path is changed based on the latest next hops from the IGP. If a better path is available, multipoint LDP attempts to switch to the better path.

Packet Forwarding

For either PIM or multipoint LDP, the device performs multicast source stream selection at the ingress interface. This preserves fabric bandwidth and maximizes forwarding performance because it:

- Avoids sending out duplicate streams across the fabric
- Prevents multiple route lookups (that result in packet drops).

For PIM, each IP multicast stream contains the same destination address. Regardless of the interface on which the packets arrive, the packets have the same route. The device checks the interface upon which each packet arrives and forwards only those that are from the primary interface. If the interface matches a backup stream interface, the device drops the packets. If the interface doesn't match either the primary or backup stream interface, the device handles the packets as exceptions in the control plane.

[Figure 137 on page 1179](#) shows this process with sample primary and backup interfaces for routers with PIM. [Figure 138 on page 1179](#) shows this similarly for switches with PIM.

Figure 137: MoFRR IP Route Lookup in the Packet Forwarding Engine on Routers

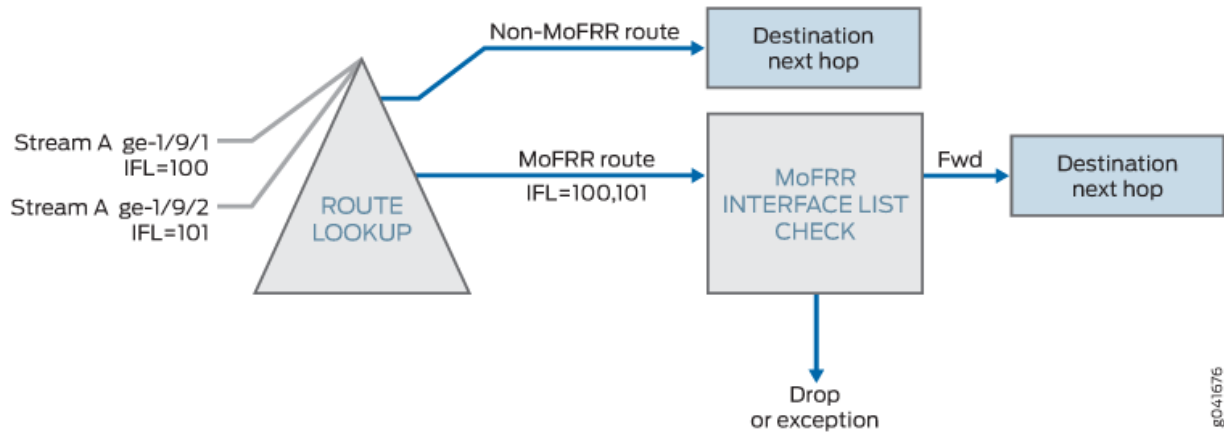
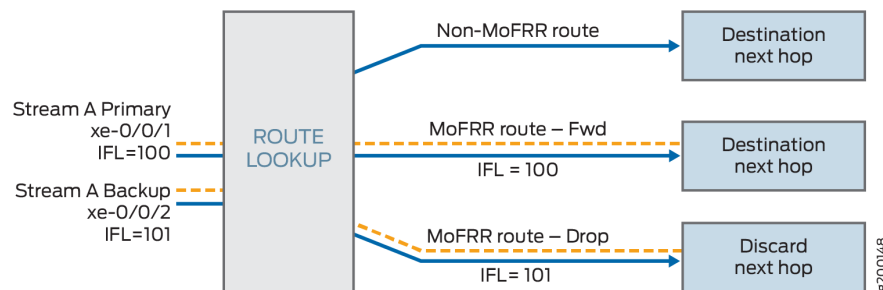


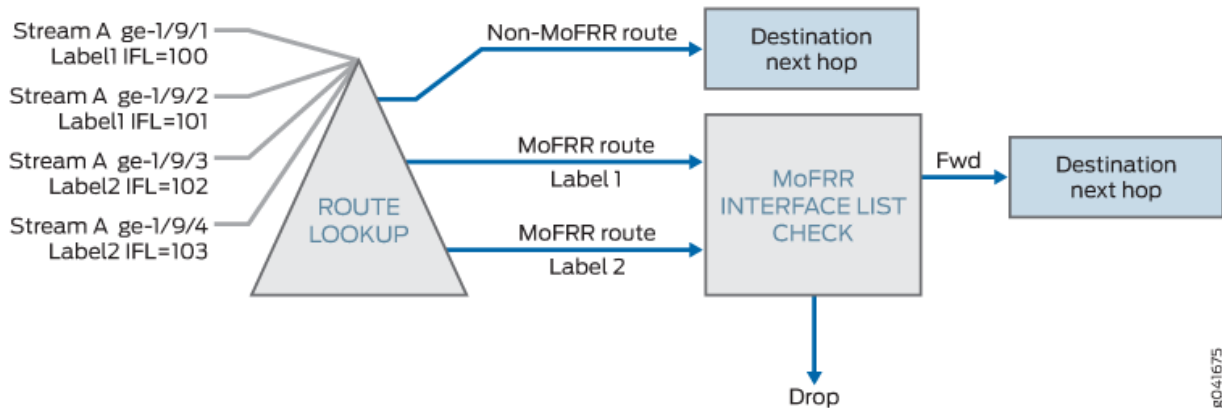
Figure 138: MoFRR IP Route Handling in the Packet Forwarding Engine on Switches



For MoFRR with multipoint LDP on routers, the device uses multiple MPLS labels to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. The device only forwards the primary label, and drops all others. Multiple interfaces can receive packets using the same label.

Figure 139 on page 1180 shows this process for routers with multipoint LDP.

Figure 139: MoFRR MPLS Route Lookup in the Packet Forwarding Engine



Limitations and Caveats

MoFRR Limitations and Caveats on Switching and Routing Devices

MoFRR has the following limitations and caveats on routing and switching devices:

- MoFRR failure detection is supported for immediate link protection of the routing device on which MoFRR is enabled and not on all the links (end-to-end) in the multicast traffic path.
- MoFRR supports fast reroute on two selected disjoint paths toward the source. Two of the selected upstream neighbors cannot be on the same interface—in other words, two upstream neighbors on a LAN segment. The same is true if the upstream interface happens to be a multicast tunnel interface.
- Detection of the maximum end-to-end disjoint upstream paths is not supported. The receiver side (egress) routing device only makes sure that there is a disjoint upstream device (the immediate previous hop). PIM and multipoint LDP do not support the equivalent of explicit route objects (EROs). Hence, disjoint upstream path detection is limited to control over the immediately previous hop device. Because of this limitation, the path to the upstream device of the previous hop selected as primary and backup might be shared.
- You might see some traffic loss in the following scenarios:
 - A better upstream path becomes available on an egress device.
 - MoFRR is enabled or disabled on the egress device while there is an active traffic stream flowing.
- PIM join load balancing for join messages for backup paths are not supported.
- For a multicast group G, MoFRR is not allowed for both (S,G) and (*,G) join messages. (S,G) join messages have precedence over (*,G).

- MoFRR is not supported for multicast traffic streams that use two different multicast groups. Each (S,G) combination is treated as a unique multicast traffic stream.
- The bidirectional PIM range is not supported with MoFRR.
- PIM dense-mode is not supported with MoFRR.
- Multicast statistics for the backup traffic stream are not maintained by PIM and therefore are not available in the operational output of `show` commands.
- Rate monitoring is not supported.

MoFRR Limitations on Switching Devices with PIM

MoFRR with PIM has the following limitations on switching devices:

- MoFRR is not supported when the upstream interface is an integrated routing and bridging (IRB) interface, which impacts other multicast features such as Internet Group Management Protocol version 3 (IGMPv3) snooping.
- Packet replication and multicast lookups while forwarding multicast traffic can cause packets to recirculate through PFEs multiple times. As a result, displayed values for multicast packet counts from the `show pfe statistics traffic` command might show higher numbers than expected in output fields such as `Input packets` and `Output packets`. You might notice this behavior more frequently in MoFRR scenarios because duplicate primary and backup streams increase the traffic flow in general.

MoFRR Limitations and Caveats on Routing Devices with Multipoint LDP

MoFRR has the following limitations and caveats on routers when used with multipoint LDP:

- MoFRR does not apply to multipoint LDP traffic received on an RSVP tunnel because the RSVP tunnel is not associated with any interface.
- Mixed upstream MoFRR is not supported. This refers to PIM multipoint LDP in-band signaling, wherein one upstream path is through multipoint LDP and the second upstream path is through PIM.
- Multipoint LDP labels as inner labels are not supported.
- If the source is reachable through multiple ingress (source-side) provider edge (PE) routing devices, multipoint LDP MoFRR is not supported.
- Targeted LDP upstream sessions are not selected as the upstream device for MoFRR.
- Multipoint LDP link protection on the backup path is not supported because there is no support for MoFRR inner labels.

Configuring Multicast-Only Fast Reroute

You can configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure.

When fast reroute is applied to unicast streams, an upstream router preestablishes MPLS label-switched paths (LSPs) or precomputes an IP loop-free alternate (LFA) fast reroute backup path to handle failure of a segment in the downstream path.

In multicast routing, the traffic distribution graphs are usually originated by the receiver. This is unlike unicast routing, which usually establishes the path from the source to the receiver. Protocols that are capable of establishing multicast distribution graphs are PIM (for IP), multipoint LDP (for MPLS) and RSVP-TE (for MPLS). Of these, PIM and multipoint LDP receivers initiate the distribution graph setup, and therefore:

- On the QFX series, MoFRR is supported in PIM domains.
- On the MX Series and SRX Series, MoFRR is supported in PIM and multipoint LDP domains.

The configuration steps are the same for enabling MoFRR for PIM on all devices that support this feature, unless otherwise indicated. Configuration steps that are not applicable to multipoint LDP MoFRR are also indicated.

(For MX Series routers only) MoFRR is supported on MX Series routers with MPC line cards. As a prerequisite, all the line cards in the router must be MPCs.

To configure MoFRR on routers or switches:

1. (For MX Series and SRX Series routers only) Set the router to enhanced IP mode.

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

2. Enable MoFRR.

```
[edit routing-options multicast]
user@host# set stream-protection
```

3. (Optional) Configure a routing policy that filters for a restricted set of multicast streams to be affected by your MoFRR configuration.

You can apply filters that are based on source or group addresses.

For example:

```
[edit policy-options]
policy-statement mofrr-select {
  term A {
    from {
      source-address-filter 225.1.1.1/32 exact;
    }
    then {
      accept;
    }
  }
  term B {
    from {
      source-address-filter 226.0.0.0/8 orlonger;
    }
    then {
      accept;
    }
  }
  term C {
    from {
      source-address-filter 227.1.1.0/24 orlonger;
      source-address-filter 227.4.1.0/24 orlonger;
      source-address-filter 227.16.1.0/24 orlonger;
    }
    then {
      accept;
    }
  }
  term D {
    from {
      source-address-filter 227.1.1.1/32 exact
    }
    then {
      reject; #MoFRR disabled
    }
  }
  ...
}
```

4. (Optional) If you configured a routing policy to filter the set of multicast groups to be affected by your MoFRR configuration, apply the policy for MoFRR stream protection.

```
[edit routing-options multicast stream-protection]
user@host# set policy policy-name
```

For example:

```
routing-options {
  multicast {
    stream-protection {
      policy mofrr-select
    }
  }
}
```

5. (Optional) In a PIM domain with MoFRR, allow MoFRR to be applied to any-source multicast (ASM) (*,G) joins.

This is not supported for multipoint LDP MoFRR.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-asm-starg
```

6. (Optional) In a PIM domain with MoFRR, allow only a disjoint RPF (an RPF on a separate plane) to be selected as the backup RPF path.

This is not supported for multipoint LDP MoFRR. In a multipoint LDP MoFRR domain, the same label is shared between parallel links to the same upstream neighbor. This is not the case in a PIM domain, where each link forms a neighbor. The `mofrr-disjoint-upstream-only` statement does not allow a backup RPF path to be selected if the path goes to the same upstream neighbor as that of the primary RPF path. This ensures that MoFRR is triggered only on a topology that has multiple RPF upstream neighbors.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-disjoint-upstream-only
```

7. (Optional) In a PIM domain with MoFRR, prevent sending join messages on the backup path, but retain all other MoFRR functionality.

This is not supported for multipoint LDP MoFRR.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-no-backup-join
```

8. (Optional) In a PIM domain with MoFRR, allow new primary path selection to be based on the unicast gateway selection for the unicast route to the source and to change when there is a change in the unicast selection, rather than having the backup path be promoted as primary. This ensures that the primary RPF hop is always on the best path.

When you include the `mofrr-primary-selection-by-routing` statement, the backup path is not guaranteed to get promoted to be the new primary path when the primary path goes down.

This is not supported for multipoint LDP MoFRR.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-primary-path-selection-by-routing
```

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

IN THIS SECTION

- [Requirements | 1186](#)
- [Overview | 1186](#)
- [CLI Quick Configuration | 1188](#)
- [Step-by-Step Configuration | 1190](#)
- [Verification | 1194](#)

This example shows how to configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure. It works by enhancing the multicast routing protocol, Protocol Independent Multicast (PIM).

MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. Data packets are received from both the primary path and the backup paths. The redundant packets are discarded at topology merge points, based on priority (weights assigned to primary and backup paths).

When a failure is detected on the primary path, the repair is made by changing the interface on which packets are accepted to the secondary interface. Because the repair is local, it is fast—greatly improving convergence times in the event of a link failure on the primary path.

Requirements

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

In this example, only the egress provider edge (PE) router has MoFRR enabled, MoFRR in a PIM domain can be enabled on any of the routers.

MoFRR is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to [network-services](#) enhanced-ip mode, and all the line-cards in the platform must be MPCs.

This example requires Junos OS Release 14.1 or later on the egress PE router.

Overview

IN THIS SECTION

- [Topology](#) | 1187

In this example, Device R3 is the egress edge router. MoFRR is enabled on this device only.

OSPF or IS-IS is used for connectivity, though any interior gateway protocol (IGP) or static routes can be used.

PIM sparse mode version 2 is enabled on all devices in the PIM domain. Device R1 serves as the rendezvous point (RP).

Device R3, in addition to MoFRR, also has PIM join load balancing enabled.

For testing purposes, routers are used to simulate the source and the receiver. Device R3 is configured to statically join the desired group by using the `set protocols igmp interface fe-1/2/15.0 static group 225.1.1.1` command. It is just joining, not listening. The fe-1/2/15.0 interface is the Device R3 interface facing the receiver. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the receiver, to make it listen to the multicast group address, this example uses `set protocols sap listen 225.1.1.1`. To make the source send multicast traffic, a multicast ping is issued from the source router. The ping command is `ping 225.1.1.1 bypass-routing interface fe-1/2/10.0 ttl 10 count 1000000000`. The fe-1/2/10.0 interface is the source interface facing Device R1.

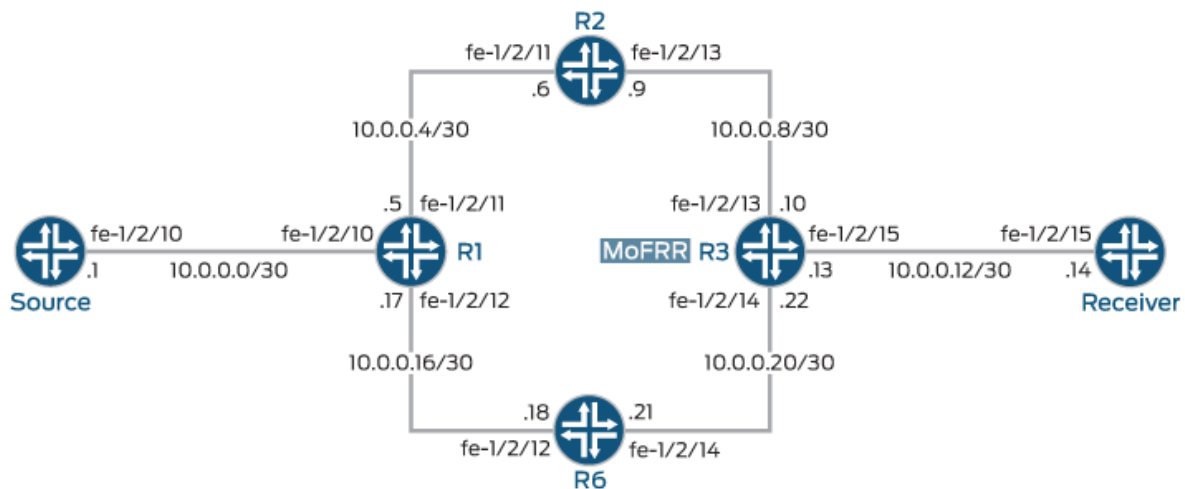
MoFRR configuration includes multiple options that are not shown in this example, but are explained separately. The options are as follows:

```
stream-protection {
  mofrr-asm-starg;
  mofrr-disjoint-upstream-only;
  mofrr-no-backup-join;
  mofrr-primary-path-selection-by-routing;
  policy policy-name;
}
```

Topology

Figure 140 on page 1187 shows the sample network.

Figure 140: MoFRR in a PIM Domain



"CLI Quick Configuration" on page 1188 shows the configuration for all of the devices in Figure 140 on page 1187.

The section "Step-by-Step Configuration" on page 1190 describes the steps on Device R3.

CLI Quick Configuration

IN THIS SECTION

- [CLI Quick Configuration | 1188](#)

CLI Quick Configuration

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

Device R1

```
set interfaces fe-1/2/10 unit 0 family inet address 10.0.0.2/30
set interfaces fe-1/2/11 unit 0 family inet address 10.0.0.5/30
set interfaces fe-1/2/12 unit 0 family inet address 10.0.0.17/30
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols ospf area 0.0.0.0 interface fe-1/2/10.0
set protocols ospf area 0.0.0.0 interface fe-1/2/11.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/12.0
set protocols pim rp local family inet address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2
```

Device R2

```
set interfaces fe-1/2/11 unit 0 family inet address 10.0.0.6/30
set interfaces fe-1/2/13 unit 0 family inet address 10.0.0.9/30
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set protocols ospf area 0.0.0.0 interface fe-1/2/11.0
set protocols ospf area 0.0.0.0 interface fe-1/2/13.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2
```

Device R3

```

set chassis network-services enhanced-ip
set interfaces fe-1/2/13 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/15 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/14 unit 0 family inet address 10.0.0.22/30
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set protocols igmp interface fe-1/2/15.0 static group 225.1.1.1
set protocols ospf area 0.0.0.0 interface fe-1/2/13.0
set protocols ospf area 0.0.0.0 interface fe-1/2/15.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/14.0
set protocols pim rp static address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim join-load-balance automatic
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-options forwarding-table export load-balancing-policy
set routing-options multicast stream-protection

```

Device R6

```

set interfaces fe-1/2/12 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/14 unit 0 family inet address 10.0.0.21/30
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set protocols ospf area 0.0.0.0 interface fe-1/2/12.0
set protocols ospf area 0.0.0.0 interface fe-1/2/14.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2

```

Device Source

```

set interfaces fe-1/2/10 unit 0 family inet address 10.0.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set protocols ospf area 0.0.0.0 interface fe-1/2/10.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Device Receiver

```
set interfaces fe-1/2/15 unit 0 family inet address 10.0.0.14/30
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set protocols sap listen 225.1.1.1
set protocols ospf area 0.0.0.0 interface fe-1/2/15.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Step-by-Step Configuration

IN THIS SECTION

- [Procedure | 1190](#)

Procedure

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 [Junos OS CLI User Guide](#).

To configure Device R3:

1. Enable enhanced IP mode.

```
[edit chassis]
user@R3# set network-services enhanced-ip
```

2. Configure the device interfaces.

```
[edit interfaces]
user@R3# set fe-1/2/13 unit 0 family inet address 10.0.0.10/30
user@R3# set fe-1/2/15 unit 0 family inet address 10.0.0.13/30
user@R3# set fe-1/2/14 unit 0 family inet address 10.0.0.22/30
user@R3# set lo0 unit 0 family inet address 192.168.0.3/32
```


3. For testing purposes only, on the interface facing Device Receiver, simulate IGMP joins.

If your test environment has receiver hosts, this step is not necessary.

```
[edit protocols igmp interface fe-1/2/15.0]
user@R3# set static group 225.1.1.1
```

4. Configure an IGP or static routes.

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface fe-1/2/13.0
user@R3# set interface fe-1/2/15.0
user@R3# set interface lo0.0 passive
user@R3# set interface fe-1/2/14.0
```

5. Configure PIM.

```
[edit protocols pim]
user@R3# set rp static address 192.168.0.1
user@R3# set interface all mode sparse
user@R3# set interface all version 2
```

6. (Optional) Configure PIM join load balancing.

```
[edit protocols pim]
user@R3# set join-load-balance automatic
```

7. (Optional) Configure per-packet load balancing.

```
[edit policy-options policy-statement load-balancing-policy]
user@R3# set then load-balance per-packet
[edit routing-options forwarding-table]
user@R3# set export load-balancing-policy
```

8. Enable MoFRR.

```
[edit routing-options multicast]
user@R3# set stream-protection
```

Results

From configuration mode, confirm your configuration by entering the show chassis, show interfaces, show protocols, show policy-options, 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@R3# show chassis
network-services enhanced-ip;
```

```
user@R3# show interfaces
fe-1/2/13 {
  unit 0 {
    family inet {
      address 10.0.0.10/30;
    }
  }
}
fe-1/2/14 {
  unit 0 {
    family inet {
      address 10.0.0.22/30;
    }
  }
}
fe-1/2/15 {
  unit 0 {
    family inet {
      address 10.0.0.13/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
```

```
        address 192.168.0.3/32;
    }
}
}
```

```
user@R3# show protocols
igmp {
  interface fe-1/2/15.0 {
    static {
      group 225.1.1.1;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/13.0;
    interface fe-1/2/15.0;
    interface lo0.0 {
      passive;
    }
    interface fe-1/2/14.0;
  }
}
pim {
  rp {
    static {
      address 192.168.0.1;
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  join-load-balance {
    automatic;
  }
}
```

```
user@R3# show policy-options
policy-statement load-balancing-policy {
```

```
    then {  
        load-balance per-packet;  
    }  
}
```

```
user@R3# show routing-options  
forwarding-table {  
    export load-balancing-policy;  
}  
multicast {  
    stream-protection;  
}
```

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

Verification

IN THIS SECTION

- [Sending Multicast Traffic Into the PIM Domain | 1194](#)
- [Verifying the Upstream Interfaces | 1195](#)
- [Checking the Multicast Routes | 1196](#)

Confirm that the configuration is working properly.

Sending Multicast Traffic Into the PIM Domain

Purpose

Use a multicast ping command to simulate multicast traffic.

Action

```
user@Source> ping 225.1.1.1 bypass-routing interface fe-1/2/10.0 ttl 10 count 1000000000  
  
PING 225.1.1.1 (225.1.1.1): 56 data bytes  
64 bytes from 10.0.0.14: icmp_seq=1 ttl=61 time=0.845 ms
```

```
64 bytes from 10.0.0.14: icmp_seq=2 ttl=61 time=0.661 ms
64 bytes from 10.0.0.14: icmp_seq=3 ttl=61 time=0.615 ms
64 bytes from 10.0.0.14: icmp_seq=4 ttl=61 time=0.640 ms
```

Meaning

The interface on Device Source, facing Device R1, is fe-1/2/10.0. Keep in mind that multicast pings have a TTL of 1 by default, so you must use the ttl option.

Verifying the Upstream Interfaces

Purpose

Make sure that the egress device has two upstream interfaces for the multicast group join.

Action

```
user@R3> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
  Source: 10.0.0.1
  Flags: sparse,spt
  Active upstream interface: fe-1/2/13.0
  Active upstream neighbor: 10.0.0.9
  MoFRR Backup upstream interface: fe-1/2/14.0
  MoFRR Backup upstream neighbor: 10.0.0.21
  Upstream state: Join to Source, No Prune to RP
  Keepalive timeout: 354
  Uptime: 00:00:06
  Downstream neighbors:
    Interface: fe-1/2/15.0
      10.0.0.13 State: Join Flags: S  Timeout: Infinity
      Uptime: 00:00:06 Time since last Join: 00:00:06
  Number of downstream interfaces: 1
```

Meaning

The output shows an active upstream interface and neighbor, and also an MoFRR backup upstream interface and neighbor.

Checking the Multicast Routes

Purpose

Examine the IP multicast forwarding table to make sure that there is an upstream RPF interface list, with a primary and a backup interface.

Action

```
user@R3> show multicast route extensive

Instance: master Family: INET

Group: 225.1.1.1
  Source: 10.0.0.1/32
  Upstream rpf interface list:
    fe-1/2/13.0 (P) fe-1/2/14.0 (B)
  Downstream interface list:
    fe-1/2/15.0
  Session description: Unknown
  Forwarding statistics are not available
  RPF Next-hop ID: 836
  Next-hop ID: 1048585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 171 seconds
  Wrong incoming interface notifications: 0
  Uptime: 00:03:09
```

Meaning

The output shows an upstream RPF interface list, with a primary and a backup interface.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Configuring Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

IN THIS SECTION

- [Requirements | 1197](#)
- [Overview | 1198](#)
- [CLI Quick Configuration | 1199](#)
- [Step-by-Step Configuration | 1201](#)
- [Verification | 1205](#)

This example shows how to configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure. It works by enhancing the multicast routing protocol, Protocol Independent Multicast (PIM).

MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. Data packets are received from both the primary path and the backup paths. The redundant packets are discarded at topology merge points, based on priority (weights assigned to primary and backup paths). When a failure is detected on the primary path, the repair is made by changing the interface on which packets are accepted to the secondary interface. Because the repair is local, it is fast—greatly improving convergence times in the event of a link failure on the primary path.

Requirements

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

This example uses QFX Series switches, and only the egress provider edge (PE) device has MoFRR enabled. This topology might alternatively include MX Series routers for the other devices where MoFRR is not enabled; in that case, substitute the corresponding interfaces for MX Series device ports used for the primary or backup multicast traffic streams.

This example requires Junos OS Release 17.4R1 or later on the device running MoFRR.

Overview

IN THIS SECTION

- [Topology | 1199](#)

In this example, Device R3 is the egress edge device. MoFRR is enabled on this device only.

OSPF or IS-IS is used for connectivity, though any interior gateway protocol (IGP) or static routes can be used.

PIM sparse mode version 2 is enabled on all devices in the PIM domain. Device R1 serves as the rendezvous point (RP).

Device R3, in addition to MoFRR, also has PIM join load balancing enabled.

For testing purposes, routing or switching devices are used to simulate the multicast source and the receiver. Device R3 is configured to statically join the desired group by using the `set protocols igmp interface xe-0/0/15.0 static group 225.1.1.1` command. It is just joining, not listening. The `xe-0/0/15.0` interface is the Device R3 interface facing the receiver. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the receiver, to listen to the multicast group address, this example uses `set protocols sap listen 225.1.1.1`. For the source to send multicast traffic, a multicast ping is issued from the source device. The ping command is `ping 225.1.1.1 bypass-routing interface xe-0/0/10.0 ttl 10 count 1000000000`. The `xe-0/0/10.0` interface is the source interface facing Device R1.


MoFRR configuration includes multiple options that are not shown in this example, but are explained separately. The options are as follows:

```
stream-protection {
  mofrr-asm-starg;
  mofrr-disjoint-upstream-only;
  mofrr-no-backup-join;
  mofrr-primary-path-selection-by-routing;
  policy policy-name;
}
```


Topology

Figure 141 on page 1199 shows the sample network.

Figure 141: MoFRR in a PIM Domain



"CLI Quick Configuration" on page 1199 shows the configuration for all of the devices in Figure 141 on page 1199.

The section "Step-by-Step Configuration" on page 1201 describes the steps on Device R3.

CLI Quick Configuration

IN THIS SECTION

- CLI Quick Configuration | 1199

CLI Quick Configuration

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

Device R1

```
set interfaces xe-0/0/10 unit 0 family inet address 10.0.0.2/30
set interfaces xe-0/0/11 unit 0 family inet address 10.0.0.5/30
```

```

set interfaces xe-0/0/12 unit 0 family inet address 10.0.0.17/30
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols ospf area 0.0.0.0 interface xe-0/0/10.0
set protocols ospf area 0.0.0.0 interface xe-0/0/11.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/12.0
set protocols pim rp local family inet address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2

```

Device R2

```

set interfaces xe-0/0/11 unit 0 family inet address 10.0.0.6/30
set interfaces xe-0/0/13 unit 0 family inet address 10.0.0.9/30
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set protocols ospf area 0.0.0.0 interface xe-0/0/11.0
set protocols ospf area 0.0.0.0 interface xe-0/0/13.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2

```

Device R3

```

set interfaces xe-0/0/13 unit 0 family inet address 10.0.0.10/30
set interfaces xe-0/0/15 unit 0 family inet address 10.0.0.13/30
set interfaces xe-0/0/14 unit 0 family inet address 10.0.0.22/30
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set protocols igmp interface xe-0/0/15.0 static group 225.1.1.1
set protocols ospf area 0.0.0.0 interface xe-0/0/13.0
set protocols ospf area 0.0.0.0 interface xe-0/0/15.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/14.0
set protocols pim rp static address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim join-load-balance automatic
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-options forwarding-table export load-balancing-policy
set routing-options multicast stream-protection

```

Device R6

```
set interfaces xe-0/0/12 unit 0 family inet address 10.0.0.18/30
set interfaces xe-0/0/14 unit 0 family inet address 10.0.0.21/30
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set protocols ospf area 0.0.0.0 interface xe-0/0/12.0
set protocols ospf area 0.0.0.0 interface xe-0/0/14.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 192.168.0.1
set protocols pim interface all mode sparse
set protocols pim interface all version 2
```

Device Source

```
set interfaces xe-0/0/10 unit 0 family inet address 10.0.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set protocols ospf area 0.0.0.0 interface xe-0/0/10.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Device Receiver

```
set interfaces xe-0/0/15 unit 0 family inet address 10.0.0.14/30
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set protocols sap listen 225.1.1.1
set protocols ospf area 0.0.0.0 interface xe-0/0/15.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Step-by-Step Configuration

IN THIS SECTION

- [Procedure | 1202](#)

Procedure

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 [Junos OS CLI User Guide](#).

To configure Device R3:

1. Configure the device interfaces.

```
[edit interfaces]
user@R3# set xe-0/0/13 unit 0 family inet address 10.0.0.10/30
user@R3# set xe-0/0/15 unit 0 family inet address 10.0.0.13/30
user@R3# set xe-0/0/14 unit 0 family inet address 10.0.0.22/30
user@R3# set lo0 unit 0 family inet address 192.168.0.3/32
```

2. For testing purposes only, on the interface facing the device labeled Receiver, simulate IGMP joins.

If your test environment has receiver hosts, this step is not necessary.

```
[edit protocols igmp interface xe-0/0/15.0]
user@R3# set static group 225.1.1.1
```

3. Configure IGP or static routes.

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface xe-0/0/13.0
user@R3# set interface xe-0/0/15.0
user@R3# set interface lo0.0 passive
user@R3# set interface xe-0/0/14.0
```

4. Configure PIM.

```
[edit protocols pim]
user@R3# set rp static address 192.168.0.1
user@R3# set interface all mode sparse
user@R3# set interface all version 2
```

5. (Optional) Configure PIM join load balancing.

```
[edit protocols pim]
user@R3# set join-load-balance automatic
```

6. (Optional) Configure per-packet load balancing.

```
[edit policy-options policy-statement load-balancing-policy]
user@R3# set then load-balance per-packet
[edit routing-options forwarding-table]
user@R3# set export load-balancing-policy
```

7. Enable MoFRR.

```
[edit routing-options multicast]
user@R3# set stream-protection
```

Results

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

```
user@R3# show interfaces
xe-0/0/13 {
  unit 0 {
    family inet {
      address 10.0.0.10/30;
    }
  }
}
xe-0/0/14 {
  unit 0 {
    family inet {
      address 10.0.0.22/30;
    }
  }
}
```

```
xe-0/0/15 {
  unit 0 {
    family inet {
      address 10.0.0.13/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.3/32;
    }
  }
}
```

```
user@R3# show protocols
igmp {
  interface xe-0/0/15.0 {
    static {
      group 225.1.1.1;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface xe-0/0/13.0;
    interface xe-0/0/15.0;
    interface lo0.0 {
      passive;
    }
    interface xe-0/0/14.0;
  }
}
pim {
  rp {
    static {
      address 192.168.0.1;
    }
  }
  interface all {
    mode sparse;
  }
}
```

```
    version 2;
  }
  join-load-balance {
    automatic;
  }
}
```

```
user@R3# show policy-options
policy-statement load-balancing-policy {
  then {
    load-balance per-packet;
  }
}
```

```
user@R3# show routing-options
forwarding-table {
  export load-balancing-policy;
}
multicast {
  stream-protection;
}
```

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

Verification

IN THIS SECTION

- [Sending Multicast Traffic Into the PIM Domain | 1206](#)
- [Verifying the Upstream Interfaces | 1206](#)
- [Checking the Multicast Routes | 1207](#)

Confirm that the configuration is working properly.

Sending Multicast Traffic Into the PIM Domain

Purpose

Use a multicast ping command to simulate multicast traffic.

Action

```
user@Source> ping 225.1.1.1 bypass-routing interface xe-0/0/10.0 ttl 10 count 100000000

PING 225.1.1.1 (225.1.1.1): 56 data bytes
64 bytes from 10.0.0.14: icmp_seq=1 ttl=61 time=0.845 ms
64 bytes from 10.0.0.14: icmp_seq=2 ttl=61 time=0.661 ms
64 bytes from 10.0.0.14: icmp_seq=3 ttl=61 time=0.615 ms
64 bytes from 10.0.0.14: icmp_seq=4 ttl=61 time=0.640 ms
```

Meaning

The interface on Device Source, facing Device R1, is xe-0/0/10.0. Keep in mind that multicast pings have a TTL of 1 by default, so you must use the ttl option.

Verifying the Upstream Interfaces

Purpose

Make sure that the egress device has two upstream interfaces for the multicast group join.

Action

```
user@R3> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
  Source: 10.0.0.1
  Flags: sparse,spt
  Active upstream interface: xe-0/0/13.0
  Active upstream neighbor: 10.0.0.9
  MoFRR Backup upstream interface: xe-0/0/14.0
  MoFRR Backup upstream neighbor: 10.0.0.21
```



```

Upstream state: Join to Source, No Prune to RP
Keepalive timeout: 354
Uptime: 00:00:06
Downstream neighbors:
  Interface: xe-0/0/15.0
    10.0.0.13 State: Join Flags: S   Timeout: Infinity
    Uptime: 00:00:06 Time since last Join: 00:00:06
Number of downstream interfaces: 1

```

Meaning

The output shows an active upstream interface and neighbor, and also an MoFRR backup upstream interface and neighbor.

Checking the Multicast Routes

Purpose

Examine the IP multicast forwarding table to make sure that there is an upstream RPF interface list, with a primary and a backup interface.

Action

```

user@R3> show multicast route extensive

Instance: master Family: INET

Group: 225.1.1.1
  Source: 10.0.0.1/32
  Upstream rpf interface list:
    xe-0/0/13.0 (P) xe-0/0/14.0 (B)
  Downstream interface list:
    xe-0/0/15.0
  Session description: Unknown
  Forwarding statistics are not available
  RPF Next-hop ID: 836
  Next-hop ID: 1048585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 171 seconds

```

```
Wrong incoming interface notifications: 0  
Uptime: 00:03:09
```

Meaning

The output shows an upstream RPF interface list, with a primary and a backup interface.

RELATED DOCUMENTATION

Understanding Multicast-Only Fast Reroute

Configuring Multicast-Only Fast Reroute

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

IN THIS SECTION

- [Requirements | 1209](#)
- [Overview | 1209](#)
- [CLI Quick Configuration | 1210](#)
- [Configuration | 1219](#)
- [Verification | 1226](#)

This example shows how to configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure.

Multipoint LDP MoFRR is used at the egress node of an MPLS network, where the packets are forwarded to an IP network. In the case of multipoint LDP MoFRR, the two paths toward the upstream provider edge (PE) router are established for receiving two streams of MPLS packets at the label-edge router (LER). One of the streams (the primary) is accepted, and the other one (the backup) is dropped at the LER. The backup stream is accepted if the primary path fails.

Requirements

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

In a multipoint LDP domain, for MoFRR to work, only the egress PE router needs to have MoFRR enabled. The other routers do not need to support MoFRR.

MoFRR is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to [network-services](#) enhanced-ip mode, and all the line-cards in the platform must be MPCs.

This example requires Junos OS Release 14.1 or later on the egress PE router.

Overview

IN THIS SECTION

- [Topology | 1209](#)

In this example, Device R3 is the egress edge router. MoFRR is enabled on this device only.

OSPF is used for connectivity, though any interior gateway protocol (IGP) or static routes can be used.

For testing purposes, routers are used to simulate the source and the receiver. Device R4 and Device R8 are configured to statically join the desired group by using the `set protocols igmp interface interface-name static group group` command. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the receivers, to make them listen to the multicast group address, this example uses `set protocols sap listen group`.

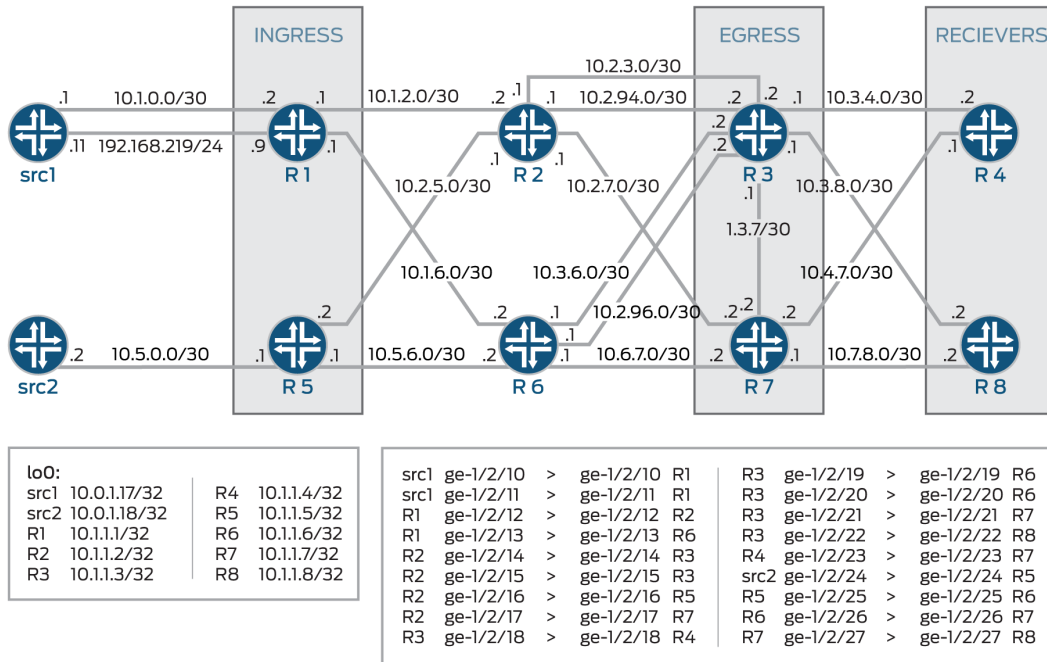
MoFRR configuration includes a policy option that is not shown in this example, but is explained separately. The option is configured as follows:

```
stream-protection {  
    policy policy-name;  
}
```

Topology

[Figure 142 on page 1210](#) shows the sample network.

Figure 142: MoFRR in a Multipoint LDP Domain



"CLI Quick Configuration" on page 1210 shows the configuration for all of the devices in Figure 142 on page 1210.

The section "Configuration" on page 1219 describes the steps on Device R3.

CLI Quick Configuration

IN THIS SECTION

- CLI Quick Configuration | 1210

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 src1

```
set interfaces ge-1/2/10 unit 0 description src1-to-R1
set interfaces ge-1/2/10 unit 0 family inet address 10.5.0.1/30
set interfaces ge-1/2/11 unit 0 description src1-to-R1
set interfaces ge-1/2/11 unit 0 family inet address 192.168.219.11/24
set interfaces lo0 unit 0 family inet address 10.0.1.17/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Device src2

```
set interfaces ge-1/2/24 unit 0 description src2-to-R5
set interfaces ge-1/2/24 unit 0 family inet address 10.5.0.2/30
set interfaces lo0 unit 0 family inet address 10.0.1.18/32
set protocols rsvp interface all
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Device R1

```
set interfaces ge-1/2/12 unit 0 description R1-to-R2
set interfaces ge-1/2/12 unit 0 family inet address 10.1.2.1/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/13 unit 0 description R1-to-R6
set interfaces ge-1/2/13 unit 0 family inet address 10.1.6.1/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/10 unit 0 description R1-to-src1
set interfaces ge-1/2/10 unit 0 family inet address 10.1.0.2/30
set interfaces ge-1/2/11 unit 0 description R1-to-src1
set interfaces ge-1/2/11 unit 0 family inet address 192.168.219.9/30
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 10.1.1.1
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 65010
set protocols bgp group ibgp neighbor 10.1.1.3
set protocols bgp group ibgp neighbor 10.1.1.7
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
```

```

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/12.0
set protocols ldp interface ge-1/2/13.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp static address 10.1.1.5
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/10.0
set protocols pim interface ge-1/2/11.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 10.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 10.1.1.7/32
orlonger
set policy-options policy-statement mldppim-ex term A from source-address-filter 10.1.0.0/30
orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 65010

```

Device R2

```

set interfaces ge-1/2/12 unit 0 description R2-to-R1
set interfaces ge-1/2/12 unit 0 family inet address 10.1.2.2/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/14 unit 0 description R2-to-R3
set interfaces ge-1/2/14 unit 0 family inet address 10.2.3.1/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces ge-1/2/16 unit 0 description R2-to-R5
set interfaces ge-1/2/16 unit 0 family inet address 10.2.5.1/30
set interfaces ge-1/2/16 unit 0 family mpls
set interfaces ge-1/2/17 unit 0 description R2-to-R7
set interfaces ge-1/2/17 unit 0 family inet address 10.2.7.1/30
set interfaces ge-1/2/17 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 description R2-to-R3
set interfaces ge-1/2/15 unit 0 family inet address 10.2.94.1/30

```

```
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set interfaces lo0 unit 0 family mpls
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 10.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 65010
```

Device R3

```
set chassis network-services enhanced-ip
set interfaces ge-1/2/14 unit 0 description R3-to-R2
set interfaces ge-1/2/14 unit 0 family inet address 10.2.3.2/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces ge-1/2/18 unit 0 description R3-to-R4
set interfaces ge-1/2/18 unit 0 family inet address 10.3.4.1/30
set interfaces ge-1/2/18 unit 0 family mpls
set interfaces ge-1/2/19 unit 0 description R3-to-R6
set interfaces ge-1/2/19 unit 0 family inet address 10.3.6.2/30
set interfaces ge-1/2/19 unit 0 family mpls
set interfaces ge-1/2/21 unit 0 description R3-to-R7
set interfaces ge-1/2/21 unit 0 family inet address 10.3.7.1/30
set interfaces ge-1/2/21 unit 0 family mpls
set interfaces ge-1/2/22 unit 0 description R3-to-R8
set interfaces ge-1/2/22 unit 0 family inet address 10.3.8.1/30
set interfaces ge-1/2/22 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 description R3-to-R2
set interfaces ge-1/2/15 unit 0 family inet address 10.2.94.2/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces ge-1/2/20 unit 0 description R3-to-R6
set interfaces ge-1/2/20 unit 0 family inet address 10.2.96.2/30
set interfaces ge-1/2/20 unit 0 family mpls
```

```

set interfaces lo0 unit 0 family inet address 10.1.1.3/32 primary
set routing-options autonomous-system 65010
set routing-options multicast stream-protection
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 10.1.1.3
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 10.1.1.1
set protocols bgp group ibgp neighbor 10.1.1.5
set protocols ospf traffic-engineering
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 p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/18.0
set protocols pim interface ge-1/2/22.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 10.1.0.1/30
orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept

```

Device R4

```

set interfaces ge-1/2/18 unit 0 description R4-to-R3
set interfaces ge-1/2/18 unit 0 family inet address 10.3.4.2/30
set interfaces ge-1/2/18 unit 0 family mpls
set interfaces ge-1/2/23 unit 0 description R4-to-R7
set interfaces ge-1/2/23 unit 0 family inet address 10.4.7.1/30
set interfaces lo0 unit 0 family inet address 10.1.1.4/32
set protocols igmp interface ge-1/2/18.0 version 3
set protocols igmp interface ge-1/2/18.0 static group 232.1.1.1 group-count 2
set protocols igmp interface ge-1/2/18.0 static group 232.1.1.1 source 192.168.219.11

```



```

set protocols igmp interface ge-1/2/18.0 static group 232.2.2.2 source 10.2.7.7
set protocols sap listen 232.1.1.1
set protocols sap listen 232.2.2.2
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface ge-1/2/23.0
set protocols pim interface ge-1/2/18.0
set protocols pim interface lo0.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 10.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 65010

```

Device R5

```

set interfaces ge-1/2/24 unit 0 description R5-to-src2
set interfaces ge-1/2/24 unit 0 family inet address 10.5.0.1/30
set interfaces ge-1/2/16 unit 0 description R5-to-R2
set interfaces ge-1/2/16 unit 0 family inet address 10.2.5.2/30
set interfaces ge-1/2/16 unit 0 family mpls
set interfaces ge-1/2/25 unit 0 description R5-to-R6
set interfaces ge-1/2/25 unit 0 family inet address 10.5.6.1/30
set interfaces ge-1/2/25 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.5/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 10.1.1.5
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 65010
set protocols bgp group ibgp neighbor 10.1.1.7
set protocols bgp group ibgp neighbor 10.1.1.3
set protocols ospf traffic-engineering

```

```
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/16.0
set protocols ldp interface ge-1/2/25.0
set protocols ldp p2mp
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/24.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 65010
```

Device R6

```
set interfaces ge-1/2/13 unit 0 description R6-to-R1
set interfaces ge-1/2/13 unit 0 family inet address 10.1.6.2/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/19 unit 0 description R6-to-R3
set interfaces ge-1/2/19 unit 0 family inet address 10.3.6.1/30
set interfaces ge-1/2/19 unit 0 family mpls
set interfaces ge-1/2/25 unit 0 description R6-to-R5
set interfaces ge-1/2/25 unit 0 family inet address 10.5.6.2/30
set interfaces ge-1/2/25 unit 0 family mpls
set interfaces ge-1/2/26 unit 0 description R6-to-R7
set interfaces ge-1/2/26 unit 0 family inet address 10.6.7.1/30
set interfaces ge-1/2/26 unit 0 family mpls
set interfaces ge-1/2/20 unit 0 description R6-to-R3
set interfaces ge-1/2/20 unit 0 family inet address 10.2.96.1/30
set interfaces ge-1/2/20 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.6/30
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp
```

Device R7

```
set interfaces ge-1/2/17 unit 0 description R7-to-R2
set interfaces ge-1/2/17 unit 0 family inet address 10.2.7.2/30
set interfaces ge-1/2/17 unit 0 family mpls
set interfaces ge-1/2/21 unit 0 description R7-to-R3
set interfaces ge-1/2/21 unit 0 family inet address 10.3.7.2/30
set interfaces ge-1/2/21 unit 0 family mpls
set interfaces ge-1/2/23 unit 0 description R7-to-R4
set interfaces ge-1/2/23 unit 0 family inet address 10.4.7.2/30
set interfaces ge-1/2/23 unit 0 family mpls
set interfaces ge-1/2/26 unit 0 description R7-to-R6
set interfaces ge-1/2/26 unit 0 family inet address 10.6.7.2/30
set interfaces ge-1/2/26 unit 0 family mpls
set interfaces ge-1/2/27 unit 0 description R7-to-R8
set interfaces ge-1/2/27 unit 0 family inet address 10.7.8.1/30
set interfaces ge-1/2/27 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.7/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 10.1.1.7
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 65010
set protocols bgp group ibgp neighbor 10.1.1.5
set protocols bgp group ibgp neighbor 10.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/17.0
set protocols ldp interface ge-1/2/21.0
set protocols ldp interface ge-1/2/26.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/27.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 10.1.0.1/30
orlonger
```

```

set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 65010
set routing-options multicast stream-protection policy mldppim-ex

```

Device R8

```

set interfaces ge-1/2/22 unit 0 description R8-to-R3
set interfaces ge-1/2/22 unit 0 family inet address 10.3.8.2/30
set interfaces ge-1/2/22 unit 0 family mpls
set interfaces ge-1/2/27 unit 0 description R8-to-R7
set interfaces ge-1/2/27 unit 0 family inet address 10.7.8.2/30
set interfaces ge-1/2/27 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.1.1.8/32
set protocols igmp interface ge-1/2/22.0 version 3
set protocols igmp interface ge-1/2/22.0 static group 232.1.1.1 group-count 2
set protocols igmp interface ge-1/2/22.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface ge-1/2/22.0 static group 232.2.2.2 source 10.2.7.7
set protocols sap listen 232.1.1.1
set protocols sap listen 232.2.2.2
set protocols rsvp interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface ge-1/2/27.0
set protocols pim interface ge-1/2/22.0
set protocols pim interface lo0.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 10.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 65010

```

Configuration

IN THIS SECTION

- Procedure | 1219

Procedure

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 [Junos OS CLI User Guide](#).

To configure Device R3:

1. Enable enhanced IP mode.

```
[edit chassis]
user@R3# set network-services enhanced-ip
```

2. Configure the device interfaces.

```
[edit interfaces]
user@R3# set ge-1/2/14 unit 0 description R3-to-R2
user@R3# set ge-1/2/14 unit 0 family inet address 10.2.3.2/30
user@R3# set ge-1/2/14 unit 0 family mpls
user@R3# set ge-1/2/18 unit 0 description R3-to-R4
user@R3# set ge-1/2/18 unit 0 family inet address 10.3.4.1/30
user@R3# set ge-1/2/18 unit 0 family mpls
user@R3# set ge-1/2/19 unit 0 description R3-to-R6
user@R3# set ge-1/2/19 unit 0 family inet address 10.3.6.2/30
user@R3# set ge-1/2/19 unit 0 family mpls
user@R3# set ge-1/2/21 unit 0 description R3-to-R7
user@R3# set ge-1/2/21 unit 0 family inet address 10.3.7.1/30
user@R3# set ge-1/2/21 unit 0 family mpls
user@R3# set ge-1/2/22 unit 0 description R3-to-R8
user@R3# set ge-1/2/22 unit 0 family inet address 10.3.8.1/30
user@R3# set ge-1/2/22 unit 0 family mpls
```

```

user@R3# set ge-1/2/15 unit 0 description R3-to-R2
user@R3# set ge-1/2/15 unit 0 family inet address 10.2.94.2/30
user@R3# set ge-1/2/15 unit 0 family mpls
user@R3# set ge-1/2/20 unit 0 description R3-to-R6
user@R3# set ge-1/2/20 unit 0 family inet address 10.2.96.2/30
user@R3# set ge-1/2/20 unit 0 family mpls
user@R3# set lo0 unit 0 family inet address 10.1.1.3/32 primary

```

3. Configure the autonomous system (AS) number.

```

user@R3# set routing-options autonomous-system 6510

```

4. Configure the routing policies.

```

[edit policy-options policy-statement mldppim-ex]
user@R3# set term B from source-address-filter 192.168.0.0/24 orlonger
user@R3# set term B from source-address-filter 192.168.219.11/32 orlonger
user@R3# set term B then accept
user@R3# set term A from source-address-filter 10.1.0.1/30 orlonger
user@R3# set term A then accept
[edit policy-options policy-statement static-route-tobgp]
user@R3# set term static from protocol static
user@R3# set term static from protocol direct
user@R3# set term static then accept

```

5. Configure PIM.

```

[edit protocols pim]
user@R3# set mldp-inband-signalling policy mldppim-ex
user@R3# set interface lo0.0
user@R3# set interface ge-1/2/18.0
user@R3# set interface ge-1/2/22.0

```

6. Configure LDP.

```

[edit protocols ldp]
user@R3# set interface all
user@R3# set p2mp

```

7. Configure an IGP or static routes.

```
[edit protocols ospf]
user@R3# set traffic-engineering
user@R3# set area 0.0.0.0 interface all
user@R3# set area 0.0.0.0 interface fxp0.0 disable
user@R3# set area 0.0.0.0 interface lo0.0 passive
```

8. Configure internal BGP.

```
[edit protocols bgp group ibgp]
user@R3# set local-address 10.1.1.3
user@R3# set peer-as 65010
user@R3# set neighbor 10.1.1.1
user@R3# set neighbor 10.1.1.5
```

9. Configure MPLS and, optionally, RSVP.

```
[edit protocols mpls]
user@R3# set interface all
[edit protocols rsvp]
user@R3# set interface all
```

10. Enable MoFRR.

```
[edit routing-options multicast]
user@R3# set stream-protection
```

Results

From configuration mode, confirm your configuration by entering the `show chassis`, `show interfaces`, `show protocols`, `show policy-options`, 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@R3# show chassis
network-services enhanced-ip;
```

```
user@R3# show interfaces
ge-1/2/14 {
  unit 0 {
    description R3-to-R2;
    family inet {
      address 10.2.3.2/30;
    }
    family mpls;
  }
}
ge-1/2/18 {
  unit 0 {
    description R3-to-R4;
    family inet {
      address 10.3.4.1/30;
    }
    family mpls;
  }
}
ge-1/2/19 {
  unit 0 {
    description R3-to-R6;
    family inet {
      address 10.3.6.2/30;
    }
    family mpls;
  }
}
ge-1/2/21 {
  unit 0 {
    description R3-to-R7;
    family inet {
```



```
        address 10.3.7.1/30;
    }
    family mpls;
}
}
ge-1/2/22 {
    unit 0 {
        description R3-to-R8;
        family inet {
            address 10.3.8.1/30;
        }
        family mpls;
    }
}
ge-1/2/15 {
    unit 0 {
        description R3-to-R2;
        family inet {
            address 10.2.94.2/30;
        }
        family mpls;
    }
}
ge-1/2/20 {
    unit 0 {
        description R3-to-R6;
        family inet {
            address 10.2.96.2/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.15.1/32;
            address 10.1.1.3/32 {
                primary;
            }
        }
    }
}
```

```
}  
}
```

```
user@R3# show protocols  
rsvp {  
    interface all;  
}  
mpls {  
    interface all;  
}  
bgp {  
    group ibgp {  
        local-address 10.1.1.3;  
        peer-as 65010;  
        neighbor 10.1.1.1;  
        neighbor 10.1.1.5;  
    }  
}  
ospf {  
    traffic-engineering;  
    area 0.0.0.0 {  
        interface all;  
        interface fxp0.0 {  
            disable;  
        }  
        interface lo0.0 {  
            passive;  
        }  
    }  
}  
ldp {  
    interface all;  
    p2mp;  
}  
pim {  
    mldp-inband-signalling {  
        policy mldppim-ex;  
    }  
    interface lo0.0;  
    interface ge-1/2/18.0;
```

```
interface ge-1/2/22.0;  
}
```

```
user@R3# show policy-options  
policy-statement mldppim-ex {  
  term B {  
    from {  
      source-address-filter 192.168.0.0/24 orlonger;  
      source-address-filter 192.168.219.11/32 orlonger;  
    }  
    then accept;  
  }  
  term A {  
    from {  
      source-address-filter 10.1.0.1/30 orlonger;  
    }  
    then accept;  
  }  
}  
policy-statement static-route-tobgp {  
  term static {  
    from protocol [ static direct ];  
    then accept;  
  }  
}
```

```
user@R3# show routing-options  
autonomous-system 65010;  
multicast {  
  stream-protection;  
}
```

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

Verification

IN THIS SECTION

- [Checking the LDP Point-to-Multipoint Forwarding Equivalency Classes | 1226](#)
- [Examining the Label Information | 1227](#)
- [Checking the Multicast Routes | 1229](#)
- [Checking the LDP Point-to-Multipoint Traffic Statistics | 1230](#)

Confirm that the configuration is working properly.

Checking the LDP Point-to-Multipoint Forwarding Equivalency Classes

Purpose

Make sure the MoFRR is enabled, and determine what labels are being used.

Action

```
user@R3> show ldp p2mp fec
```

```
LDP P2MP FECs:
```

```
P2MP root-addr 10.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
```

```
MoFRR enabled
```

```
Fec type: Egress (Active)
```

```
Label: 301568
```

```
P2MP root-addr 10.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
```

```
MoFRR enabled
```

```
Fec type: Egress (Active)
```

```
Label: 301600
```

Meaning

The output shows that MoFRR is enabled, and it shows that the labels 301568 and 301600 are being used for the two multipoint LDP point-to-multipoint LSPs.

Examining the Label Information

Purpose

Make sure that the egress device has two upstream interfaces for the multicast group join.

Action

```

user@R3> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Address: 0x2735208
        Next-hop reference count: 3
        Next hop type: Router, Next hop index: 1397
        Address: 0x2735d2c
        Next-hop reference count: 3
        Next hop: 10.3.8.2 via ge-1/2/22.0
        Label operation: Pop
        Load balance label: None;
        Next hop type: Router, Next hop index: 1395
        Address: 0x2736290
        Next-hop reference count: 3
        Next hop: 10.3.4.2 via ge-1/2/18.0
        Label operation: Pop
        Load balance label: None;
        State: <Active Int AckRequest MulticastRPF>
        Local AS: 65010
        Age: 54:05      Metric: 1
        Validation State: unverified
        Task: LDP
        Announcement bits (1): 0-KRT
        AS path: I
        FECs bound to route: P2MP root-addr 10.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
Primary Upstream : 10.1.1.3:0--10.1.1.2:0
      RPF Nexthops :
        ge-1/2/15.0, 10.2.94.1, Label: 301568, weight: 0x1
        ge-1/2/14.0, 10.2.3.1, Label: 301568, weight: 0x1
Backup Upstream : 10.1.1.3:0--10.1.1.6:0
      RPF Nexthops :

```

```

ge-1/2/20.0, 10.2.96.1, Label: 301584, weight: 0xffff
ge-1/2/19.0, 10.3.6.1, Label: 301584, weight: 0xffff

```

```
user@R3> show route label 301600 detail
```

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
```

```
301600 (1 entry, 1 announced)
```

```

*LDP Preference: 9
Next hop type: Flood
Address: 0x27356b4
Next-hop reference count: 3
Next hop type: Router, Next hop index: 1520
Address: 0x27350f4
Next-hop reference count: 3
Next hop: 10.3.8.2 via ge-1/2/22.0
Label operation: Pop
Load balance label: None;
Next hop type: Router, Next hop index: 1481
Address: 0x273645c
Next-hop reference count: 3
Next hop: 10.3.4.2 via ge-1/2/18.0
Label operation: Pop
Load balance label: None;
State: <Active Int AckRequest MulticastRPF>
Local AS: 65010
Age: 54:25 Metric: 1
Validation State: unverified
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
Primary Upstream : 10.1.1.3:0--10.1.1.6:0
RPF Nexthops :
    ge-1/2/20.0, 10.2.96.1, Label: 301600, weight: 0x1
    ge-1/2/19.0, 10.3.6.1, Label: 301600, weight: 0x1
Backup Upstream : 10.1.1.3:0--1.1.1.2:0
RPF Nexthops :
    ge-1/2/15.0, 10.2.94.1, Label: 301616, weight: 0xffff
    ge-1/2/14.0, 10.2.3.1, Label: 301616, weight: 0xffff

```

Meaning

The output shows the primary upstream paths and the backup upstream paths. It also shows the RPF next hops.

Checking the Multicast Routes

Purpose

Examine the IP multicast forwarding table to make sure that there is an upstream RPF interface list, with a primary and a backup interface.

Action

```

user@R3> show ldp p2mp path
P2MP path type: Transit/Egress
  Output Session (label): 10.1.1.2:0 (301568) (Primary)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 10.2.94.1, 301568, 1
                   Interface ge-1/2/20.0, 10.2.96.1, 301584, 65534
                   Interface ge-1/2/14.0, 10.2.3.1, 301568, 1
                   Interface ge-1/2/19.0, 10.3.6.1, 301584, 65534
  Attached FECs:  P2MP root-addr 10.1.1.1, grp: 232.1.1.1, src: 192.168.219.11 (Active)
P2MP path type: Transit/Egress
  Output Session (label): 10.1.1.6:0 (301584) (Backup)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 10.2.94.1, 301568, 1
                   Interface ge-1/2/20.0, 10.2.96.1, 301584, 65534
                   Interface ge-1/2/14.0, 10.2.3.1, 301568, 1
                   Interface ge-1/2/19.0, 10.3.6.1, 301584, 65534
  Attached FECs:  P2MP root-addr 10.1.1.1, grp: 232.1.1.1, src: 192.168.219.11 (Active)
P2MP path type: Transit/Egress
  Output Session (label): 10.1.1.6:0 (301600) (Primary)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 10.2.94.1, 301616, 65534
                   Interface ge-1/2/20.0, 10.2.96.1, 301600, 1
                   Interface ge-1/2/14.0, 10.2.3.1, 301616, 65534
                   Interface ge-1/2/19.0, 10.3.6.1, 301600, 1

```

```

Attached FECs: P2MP root-addr 10.1.1.1, grp: 232.1.1.2, src: 192.168.219.11 (Active)
P2MP path type: Transit/Egress
Output Session (label): 10.1.1.2:0 (301616) (Backup)
Egress Nexthops: Interface ge-1/2/18.0
                  Interface ge-1/2/22.0
RPF Nexthops:    Interface ge-1/2/15.0, 10.2.94.1, 301616, 65534
                  Interface ge-1/2/20.0, 10.2.96.1, 301600, 1
                  Interface ge-1/2/14.0, 10.2.3.1, 301616, 65534
                  Interface ge-1/2/19.0, 10.3.6.1, 301600, 1
Attached FECs: P2MP root-addr 10.1.1.1, grp: 232.1.1.2, src: 192.168.219.11 (Active)

```

Meaning

The output shows primary and backup sessions, and RPF next hops.

Checking the LDP Point-to-Multipoint Traffic Statistics

Purpose

Make sure that both primary and backup statistics are listed.

Action

```
user@R3> show ldp traffic-statistics p2mp
```

P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes Shared
10.1.1.1:232.1.1.1,192.168.219.11, Label: 301568	10.3.8.2	0	0 No
	10.3.4.2	0	0 No
10.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route	10.3.4.2	0	0 No
	10.3.8.2	0	0 No
10.1.1.1:232.1.1.2,192.168.219.11, Label: 301600	10.3.8.2	0	0 No
	10.3.4.2	0	0 No
10.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route	10.3.4.2	0	0 No
	10.3.8.2	0	0 No

Meaning

The output shows both primary and backup routes with the labels.

Enable Multicast Between Layer 2 and Layer 3 Devices Using Snooping

IN THIS CHAPTER

- [Multicast Snooping on MX Series Routers | 1232](#)
- [Example: Configuring Multicast Snooping | 1233](#)
- [Example: Configuring Multicast Snooping for a Bridge Domain | 1245](#)
- [Configuring Multicast Snooping to Ignore Spanning Tree Topology Change Messages | 1246](#)
- [Configuring Graceful Restart for Multicast Snooping | 1248](#)
- [Multicast Snooping for VPLS | 1250](#)

Multicast Snooping on MX Series Routers

Because MX Series routers can support both Layer 3 and Layer 2 functions at the same time, you can configure the Layer 3 multicast protocols Protocol Independent Multicast (PIM) and the Internet Group Membership Protocol (IGMP) as well as Layer 2 VLANs on an MX Series router.

Normal encapsulation rules restrict Layer 2 processing to accessing information in the frame header and Layer 3 processing to accessing information in the packet header. However, in some cases, an interface running a Layer 2 protocol needs information available only at Layer 3. In multicast applications, the VLANs need the group membership information and multicast tree information available to the Layer 3 IGMP and PIM protocols. In these cases, the Layer 3 configurations can use PIM or IGMP snooping to provide the needed information at the VLAN level.

For information about configuring multicast snooping for the operational details of a Layer 3 protocol on behalf of a Layer 2 spanning-tree protocol process, see "[Understanding Multicast Snooping and VPLS Root Protection](#)" on page 1234.

Snooping configuration statements and examples are not included in the [Junos OS Layer 2 Switching and Bridging Library for Routing Devices](#). For more information about configuring PIM and IGMP snooping, see the [Junos OS Multicast Protocols User Guide](#).

RELATED DOCUMENTATION

[Understanding Multicast Snooping and VPLS Root Protection | 1234](#)

[Configuring Multicast Snooping to Ignore Spanning Tree Topology Change Messages | 1246](#)

[Example: Configuring Multicast Snooping for a Bridge Domain | 1245](#)

Example: Configuring Multicast Snooping

IN THIS SECTION

- [Understanding Multicast Snooping | 1233](#)
- [Understanding Multicast Snooping and VPLS Root Protection | 1234](#)
- [Configuring Multicast Snooping | 1235](#)
- [Example: Configuring Multicast Snooping | 1236](#)
- [Enabling Bulk Updates for Multicast Snooping | 1243](#)
- [Enabling Multicast Snooping for Multichassis Link Aggregation Group Interfaces | 1244](#)

Understanding Multicast Snooping

Network devices such as routers operate mainly at the packet level, or Layer 3. Other network devices such as bridges or LAN switches operate mainly at the frame level, or Layer 2. Multicasting functions mainly at the packet level, Layer 3, but there is a way to map Layer 3 IP multicast group addresses to Layer 2 MAC multicast group addresses at the frame level.

Routers can handle both Layer 2 and Layer 3 addressing information because the frame and its addresses must be processed to access the encapsulated packet inside. Routers can run Layer 3 multicast protocols such as PIM or IGMP and determine where to forward multicast content or when a host on an interface joins or leaves a group. However, bridges and LAN switches, as Layer 2 devices, are not supposed to have access to the multicast information inside the packets that their frames carry.

How then are bridges and other Layer 2 devices to determine when a device on an interface joins or leaves a multicast tree, or whether a host on an attached LAN wants to receive the content of a particular multicast group?

The answer is for the Layer 2 device to implement multicast snooping. Multicast snooping is a general term and applies to the process of a Layer 2 device “snooping” at the Layer 3 packet content to determine which actions are taken to process or forward a frame. There are more specific forms of

snooping, such as IGMP snooping or PIM snooping. In all cases, snooping involves a device configured to function at Layer 2 having access to normally “forbidden” Layer 3 (packet) information. Snooping makes multicasting more efficient in these devices.

SEE ALSO

[Multicast Overview | 2](#)

Understanding Multicast Snooping and VPLS Root Protection

Snooping occurs when a Layer 2 protocol such as a spanning-tree protocol is aware of the operational details of a Layer 3 protocol such as the Internet Group Management Protocol (IGMP) or other multicast protocol. Snooping is necessary when Layer 2 devices such as VLAN switches must be aware of Layer 3 information such as the media access control (MAC) addresses of members of a multicast group.

VPLS root protection is a spanning-tree protocol process in which only one interface in a multihomed environment is actively forwarding spanning-tree protocol frames. This protects the root of the spanning tree against bridging loops, but also prevents both devices in the multihomed topology from snooped information, such as IGMP membership reports.

For example, consider a collection of multicast-capable hosts connected to two customer edge (CE) routers (CE1 and CE2) which are connected to each other (a CE1–CE2 link is configured) and multihomed to two provider edge (PE) routers (PE1 and PE2, respectively). The active PE only receives forwarded spanning-tree protocol information on the active PE–CE link, due to root protection operation. As long as the CE1–CE2 link is operational, this is not a problem. However, if the link between CE1 and CE2 fails, and the other PE becomes the active spanning-tree protocol link, no multicast snooping information is available on the new active PE. The new active PE will not forward multicast traffic to the CE and the hosts serviced by this CE router.

The service outage is corrected once the hosts send new group membership IGMP reports to the CE routers. However, the service outage can be avoided if multicast snooping information is available to both PEs in spite of normal spanning-tree protocol root protection operation.

You can configure multicast snooping to ignore messages about spanning tree topology changes on bridge domains on virtual switches and bridge domains default routing switches. You can use the `ignore-stp-topology-change` command to ignore messages about spanning tree topology changes

SEE ALSO

[Understanding VPLS Multihoming](#)

[Junos OS Layer 2 Switching and Bridging Library for Routing Devices](#)

[Multicast Snooping on MX Series Routers | 1232](#)

[Junos OS Layer 2 Switching and Bridging Library for Routing Devices](#)

[Configuring Multicast Snooping to Ignore Spanning Tree Topology Change Messages | 1246](#)

[Junos OS Layer 2 Switching and Bridging Library for Routing Devices](#)

[Example: Configuring Multicast Snooping for a Bridge Domain | 1245](#)

[Junos OS Layer 2 Switching and Bridging Library for Routing Devices](#)

[Junos OS Multicast Protocols User Guide](#)

ignore-stp-topology-change

Configuring Multicast Snooping

To configure the general multicast snooping parameters for MX Series routers, include the `multicast-snooping-options` statement:

```
multicast-snooping-options {
  flood-groups [ ip-addresses ];
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  graceful-restart <restart-duration seconds>;
  ignore-stp-topology-change;
  multichassis-lag-replicate-state;
  nexthop-hold-time milliseconds;
  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*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

By default, multicast snooping is disabled. You can enable multicast snooping in VPLS or virtual switch instance types in the instance hierarchy.

If there are multiple bridge domains configured under a VPLS or virtual switch instance, the multicast snooping options configured at the instance level apply to all the bridge domains.



NOTE: The `ignore-stp-topology-change` statement is supported for the **virtual-switch** routing instance type only and is not supported under the `[edit logical-systems]` hierarchy.



NOTE: The `nexthop-hold-time` statement is supported only at the `[edit routing-instances routing-instance-name]` hierarchy, and only for an instance type of **virtual-switch** or **vpls**.

SEE ALSO

[Configuring IGMP Snooping | 147](#)

[Example: Configuring IGMP Snooping | 149](#)

[Configuring VLAN-Specific IGMP Snooping Parameters | 148](#)

[Configuring IGMP Snooping Trace Operations | 158](#)

Example: Configuring Multicast Snooping

IN THIS SECTION

- [Requirements | 1236](#)
- [Overview and Topology | 1237](#)
- [Configuration | 1239](#)
- [Verification | 1242](#)

This example shows how to configure multicast snooping in a bridge or VPLS routing-instance scenario.

Requirements

This example uses the following hardware components:

- One MX Series router
- One Layer 3 device functioning as a multicast router

Before you begin:

- Configure the interfaces.

- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure a multicast protocol. This feature works with the following multicast protocols:
 - DVMRP
 - PIM-DM
 - PIM-SM
 - PIM-SSM

Overview and Topology

IN THIS SECTION

- [Topology](#) | 1239

IGMP snooping prevents Layer 2 devices from indiscriminately flooding multicast traffic out all interfaces. The settings that you configure for multicast snooping help manage the behavior of IGMP snooping.

You can configure multicast snooping options on the default master instance and on individual bridge or VPLS instances. The default master instance configuration is global and applies to all individual bridge or VPLS instances in the logical router. The configuration for the individual instances overrides the global configuration.

This example includes the following statements:

- **flood-groups**—Enables you to list multicast group addresses for which traffic must be flooded. This setting is useful for making sure that IGMP snooping does not prevent necessary multicast flooding. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms. For example, OSPF uses 224.0.0.5 for all OSPF routers.
- **forwarding-cache**—Specifies how forwarding entries are aged out and how the number of entries is controlled.

You can configure threshold values on the forwarding cache to suppress (suspend) snooping when the cache entries reach a certain maximum and reuse the cache when the number falls to another threshold value. By default, no threshold values are enabled on the router.

The suppress threshold suppresses new multicast forwarding cache entries. An optional reuse threshold specifies the point at which the router begins to create new multicast forwarding cache entries. The range for both thresholds is from 1 through 200,000. If configured, the reuse value must be less than the suppression value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.

- **graceful-restart**—Configures the time after which routes learned before a restart are replaced with routes relearned. If graceful restart for multicast snooping is disabled, snooping information is lost after a Routing Engine restart.

By default, the graceful restart duration is 180 seconds (3 minutes). You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.

- **ignore-stp-topology-change**—Configures the MX Series router to ignore messages about the spanning-tree topology state change.

By default the IGMP snooping process on an MX Series router detects interface state changes made by any of the spanning tree protocols (STPs).

In a VPLS multihoming environment where two PE routers are connected to two interconnected CE routers and STP root protection is enabled on the PE routers, one of the PE router interfaces is in forwarding state and the other is in blocking state.

If the link interconnecting the two CE routers fails, the PE router interface in blocking state transitions to the forwarding state.

The PE router interface does not wait to receive membership reports in response to the next general or group-specific query. Instead, the IGMP snooping process sends a general query message toward the CE router. The hosts connected to the CE router reply with reports for all groups they are interested in.

When the link interconnecting the two CE routers is restored, the original spanning-tree state on both PE routers is restored. The forwarding PE receives a spanning-tree topology change message and sends a general query message toward the CE router to immediately reconstruct the group membership state.

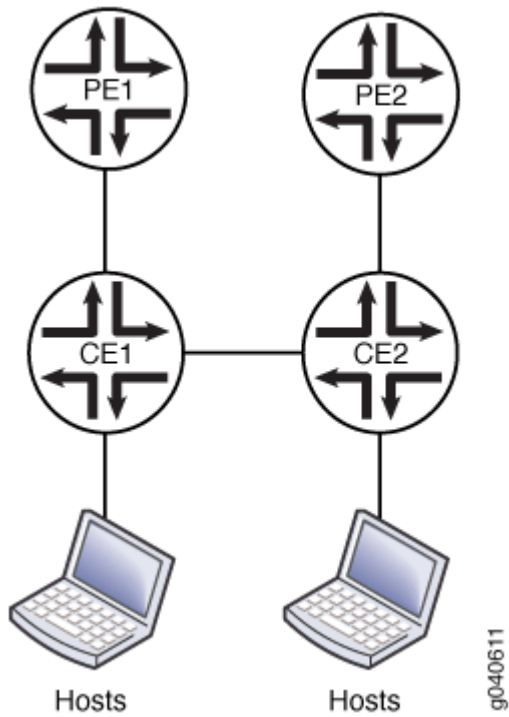


NOTE: The `ignore-stp-topology-change` statement is supported for the **virtual-switch** routing instance type only.

Topology

Figure 143 on page 1239 shows a VPLS multihoming topology in which a customer network has two CE devices with a link between them. Each CE is connected to one PE.

Figure 143: VPLS Multihoming Topology



Configuration

IN THIS SECTION

- Procedure | 1240

Procedure

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.

```
set bridge-domains domain1 multicast-snooping-options forwarding-cache threshold suppress 100
set bridge-domains domain1 multicast-snooping-options forwarding-cache threshold reuse 50
set bridge-domains domain1 multicast-snooping-options graceful-restart restart-duration 120
set routing-instances ce1 instance-type virtual-switch
set routing-instances ce1 bridge-domains domain1 domain-type bridge
set routing-instances ce1 bridge-domains domain1 vlan-id 100
set routing-instances ce1 bridge-domains domain1 interface ge-0/3/9.0
set routing-instances ce1 bridge-domains domain1 interface ge-0/0/6.0
set routing-instances ce1 bridge-domains domain1 multicast-snooping-options flood-groups
224.0.0.5
set routing-instances ce1 bridge-domains domain1 multicast-snooping-options ignore-stp-topology-
change
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the [Junos OS CLI User Guide](#).

To configure IGMP snooping:

1. Configure multicast snooping settings in the master routing instance.

```
[edit bridge-domains domain1]
user@host# set multicast-snooping-options forwarding-cache threshold suppress 100 reuse 50
user@host# set multicast-snooping-options graceful-restart 120
```

2. Configure the routing instance.

```
[edit routing-instances ce1]
user@host# set instance-type virtual-switch
```

3. Configure the bridge domain in the routing instance.

```
[edit routing-instances ce1 bridge-domains domain1]
user@host# set domain-type bridge
user@host# set interface ge-0/0/6.0
user@host# set interface ge-0/3/9.0
user@host# set vlan-id 100
```

4. Configure flood groups.

```
[edit routing-instances ce1 bridge-domains domain1]
user@host# set multicast-snooping-options flood-groups 224.0.0.5
```

5. Configure the router to ignore messages about spanning-tree topology state changes.

```
[edit routing-instances ce1 bridge-domains domain1]
user@host# set multicast-snooping-options ignore-stp-topology-change
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show bridge-domains` and `show routing-instances` commands.

```
user@host# show bridge-domains
domain1 {
  multicast-snooping-options {
    forwarding-cache {
      threshold {
        suppress 100;
        reuse 50;
      }
    }
  }
}
```

```

}
}

```

```

user@host# show routing-instances
ce1 {
  instance-type virtual-switch;
  bridge-domains {
    domain1 {
      domain-type bridge;
      vlan-id 100;
      interface ge-0/3/9.0; ## 'ge-0/3/9.0' is not defined
      interface ge-0/0/6.0; ## 'ge-0/0/6.0' is not defined
      multicast-snooping-options {
        flood-groups 224.0.0.5;
        ignore-stp-topology-change;
      }
    }
  }
}
}

```

Verification

To verify the configuration, run the following commands:

- **show igmp snooping interface**
- **show igmp snooping membership**
- **show igmp snooping statistics**
- **show multicast snooping route**
- **show route table**

SEE ALSO

[Example: Configuring IGMP Snooping | 149](#)

[Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network](#)

query-response-interval

Enabling Bulk Updates for Multicast Snooping

Whenever an individual interface joins or leaves a multicast group, a new next hop entry is installed in the routing table and the forwarding table. You can use the `nexthop-hold-time` statement to specify a time, from 1 through 1000 milliseconds (ms), during which outgoing interface changes are accumulated and then updated in bulk to the routing table and forwarding table. Bulk updating reduces the processing time and memory overhead required to process join and leave messages. This is useful for applications such as Internet Protocol television (IPTV), in which users changing channels can create thousands of interfaces joining or leaving a group in a short period. In IPTV scenarios, typically there is a relatively small and controlled number of streams and a high number of outgoing interfaces. Using bulk updates can reduce the join delay.

In this example, you configure a hold-time of 20 milliseconds for **instance-type virtual-switch**, using the `nexthop-hold-time` statement:

1. Enable the `nexthop-hold-time` statement by configuring it under **multicast-snooping-options**, using 20 milliseconds for the time value.

```
[edit routing-instances vs]
multicast-snooping-options {
  nexthop-hold-time 20;
}
```

2. Use the `show multicast snooping route` command to verify that the bulk updates feature is turned on.

```
user@host> show multicast snooping route instance vs
Nexthop Bulking: ON
Family: INET
Group: 224.0.0.0
```

You can include the `nexthop-hold-time` statement only for routing-instance types of **virtual-switch** or **vpls** at the following hierarchy level.

- [edit routing-instances *routing-instance-name* multicast-snooping-options]

If the `nexthop-hold-time` statement is deleted from the router configuration, bulk updates are disabled.

SEE ALSO

[*multicast-snooping-options*](#)

[*nexthop-hold-time*](#)

Enabling Multicast Snooping for Multichassis Link Aggregation Group Interfaces

Include the `multichassis-lag-replicate-state` statement at the `[edit multicast-snooping-options]` hierarchy level to enable IGMP snooping and state replication for multichassis link aggregation group (MC-LAG) interfaces.

```
[edit]
multicast-snooping-options {
    multichassis-lag-replicate-state;
}
```

Replicating join and leave messages between links of a dual-link MC-LAG interface enables faster recovery of membership information for MC-LAG interfaces that experience service interruption.

Without state replication, if a dual-link MC-LAG interface experiences a service interruption (for example, if an active link switches to standby), the membership information for the interface is recovered by generating an IGMP query to the network. This method can take from 1 through 10 seconds to complete, which might be too long for some applications.

When state replication is provided for MC-LAG interfaces, IGMP join or leave messages received on an MC-LAG device are replicated from the active MC-LAG link to the standby link through an Interchassis Communication Protocol (ICCP) connection. The standby link processes the messages as if they were received from the corresponding active MC-LAG link, except it does not add itself as a next hop and it does not flood the message to the network. After a failover, the multicast membership status of the link can be recovered within a few seconds or less by retrieving the replicated messages.

This example enables state replication for MC-LAG interfaces:

1. Enable state replication for MC-LAG interfaces on the routing device.

```
user@host# set multicast-snooping-options multichassis-lag-replicate-state
```

After you commit the configuration, multicast snooping automatically identifies the active link during initialization or after failover, and replicates data between the active and standby links without administrator intervention.

2. Use the `show igmp snooping interface` command to display the state for MC-LAG interfaces.

```
user@host> show igmp snooping interface

Learning-Domain: default
Interface: ae0.1
State:           Up Groups:      1
```

```

mc-lag state: standby
Immediate leave: Off
Router interface: no
Interface: ge-0/1/3.100
State:          Up Groups:    1
Immediate leave: Off
Router interface: no
Interface: ae1.2
State:          Up Groups:    1
mc-lag state: standby
Immediate leave: Off
Router interface: no

```



NOTE: You can use the `show igmp snooping membership` command to display group membership information for the links of MC-LAG interfaces.

If you delete the `multicast-lag-replicate-state` statement or the configuration of IGMP snooping, replication between MC-LAG links stops within the hierarchy level from which the configuration was deleted. Then, multicast membership is recovered as needed by generating standard IGMP queries over the network.

SEE ALSO

| [*multichassis-lag-replicate-state*](#)

Example: Configuring Multicast Snooping for a Bridge Domain

This example configures the multicast snooping option for a bridge domain named **Ignore-STP** in a virtual switch routing instance named `vs_routing_instance_multihomed_CEs`:

```

[edit]
routing-instances {
  vs_routing_instance_multihomed_CEs {
    instance-type virtual-switch;
    bridge-domains {
      bd_ignore_STP {
        multicast-snooping-options {

```

```
        ignore-stp-topology-change;
    }
}
}
```



NOTE: This is not a complete router configuration.

RELATED DOCUMENTATION

[Multicast Snooping on MX Series Routers | 1232](#)

[Understanding Multicast Snooping and VPLS Root Protection | 1234](#)

[Configuring Multicast Snooping to Ignore Spanning Tree Topology Change Messages | 1246](#)

Configuring Multicast Snooping to Ignore Spanning Tree Topology Change Messages

You can configure the multicast snooping process for a virtual switch to ignore VPLS root protection topology change messages.

Before you begin, complete the following tasks:

1. Configure the spanning-tree protocol. For configuration details, see one of the following topics:
 - [Configuring Rapid Spanning Tree Protocol](#)
 - [Configuring Multiple Spanning Tree Protocol](#)
 - [Configuring VLAN Spanning Tree Protocol](#)
2. Configure VPLS root protection. For configuration details, see one of the following topics:
 - [Configuring VPLS Root Protection Topology Change Actions to Control Individual VLAN Spanning-Tree Behavior](#)

To configure multicast snooping to ignore spanning tree topology change messages:

1. Configure a **virtual-switch** routing instance to isolate a LAN segment with its VSTP instance.
 - a.

Enable configuration of a virtual switch routing instance:

```
[edit]
user@host# edit routing-instances routing-instance-name
user@host# set instance-type virtual-switch
```

You can configure multicast snooping to ignore messages about spanning tree topology changes for the **virtual-switch** routing-instance type only.

- b. Enable configuration of a bridge domain:

```
[edit routing-instances routing-instance-name]
user@host# edit bridge-domains bridge-domain-name
user@host# set domain-type bridge
```

- c. Configure the logical interfaces for the bridge domain in the virtual switch:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
user@host# set interface interface-name
```

- d. Configure the VLAN identifiers for the bridge domain in the virtual switch. For detailed information, see [Configuring a Virtual Switch Routing Instance on MX Series Routers](#).

2. Configure the multicast snooping process to ignore any spanning tree topology change messages sent to the virtual switch routing instance:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
user@host# set multicast-snooping-options ignore-stp-topology-change
```

3. Verify the configuration of multicast snooping for the virtual-switch routing instance to ignore spanning tree topology change messages:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]
user@host# top
user@host# show routing-instances

routing-instance-name {
    instance-type virtual-switch;
    bridge-domains {
        bridge-domain-name {
```

```
domain-type bridge {  
  interface interface-name;  
  ...VLAN-identifiers-configuration...  
  multicast-snooping-options {  
    ignore-stp-topology-change;  
  }  
}  
}
```

RELATED DOCUMENTATION

[Multicast Snooping on MX Series Routers | 1232](#)

[Understanding Multicast Snooping and VPLS Root Protection | 1234](#)

[Example: Configuring Multicast Snooping for a Bridge Domain | 1245](#)

Configuring Graceful Restart for Multicast Snooping

When graceful restart is enabled for multicast snooping, no data traffic is lost during a process restart or a graceful Routing Engine switchover (GRES). Graceful restart can be configured for multicast snooping either at the global level or at the level of individual routing instances.

At the global level, graceful restart is enabled by default for multicast snooping. To change this default setting, you can configure the `disable` statement at the `[edit multicast-snooping-options graceful-restart]` hierarchy level:

```
multicast-snooping-options {  
  graceful-restart disable;  
}
```

To configure graceful restart for multicast snooping on a global level:

1. Configure the duration for graceful restart.

```
[edit multicast-snooping-options graceful-restart]  
user@host# set restart-duration 200
```

The range for `restart-duration` is from 0 through 300 seconds. The default value is 180 seconds. After this period, the Routing Engine resumes normal multicast operation.

You can also set the `graceful-restart` statement for an individual routing instance level at the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options] hierarchy level.

2. Verify your configuration by using the `show multicast-snooping-options` command.

```
[edit]
user@host# show multicast-snooping-options
```

```
graceful-restart {
  restart-duration 200;
}
```

3. Commit the configuration.

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

To configure graceful restart for multicast snooping for an individual routing instance level:

1. Configure the duration for graceful restart.

```
[edit routing-instances ri1 multicast-snooping-options graceful-restart]
user@host# set restart-duration 200
```

The range for `restart-duration` is from 0 through 300 seconds. The default value is 180 seconds. After this period, the Routing Engine resumes normal multicast operation.



NOTE: You can also set the `graceful-restart` statement for an individual routing instance level at the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options] hierarchy level.

2. Verify your configuration by using the `show routing-instances routing-instance-name multicast-snooping-options` command.

```
[edit]
user@host# show routing-instances ri1 multicast-snooping-options
```

```
graceful-restart {
  restart-duration 200;
}
```

3. Commit the configuration.

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

RELATED DOCUMENTATION

[Example: Configuring Multicast Snooping | 1233](#)

graceful-restart (Multicast Snooping)

Multicast Snooping for VPLS

IN THIS SECTION

- [Understanding PIM Snooping for VPLS | 1250](#)
- [Example: Configuring PIM Snooping for VPLS | 1252](#)
- [IGMP and MLD Snooping for VPLS | 1267](#)

Understanding PIM Snooping for VPLS

There are two ways to direct PIM control packets:

- By the use of PIM snooping
- By the use of PIM proxying

PIM snooping configures a device to examine and operate only on PIM hello and join/prune packets. A PIM snooping device snoops PIM hello and join/prune packets on each interface to find interested multicast receivers and populates the multicast forwarding tree with this information. PIM snooping differs from PIM proxying in that both PIM hello and join/prune packets are transparently flooded in the VPLS as opposed to the flooding of only hello packets in the case of PIM proxying. PIM snooping is configured on PE routers connected through pseudowires. PIM snooping ensures that no new PIM packets are generated in the VPLS, with the exception of PIM messages sent through LDP on pseudowires.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

A device that supports PIM snooping snoops hello packets received on attachment circuits. It does not introduce latency in the VPLS core when it forwards PIM join/prune packets.

To configure PIM snooping on a PE router, use the `pim-snooping` statement at the `[edit routing-instances instance-name protocols]` hierarchy level:

```
routing-instances {
  customer {
    instance-type vpls;
    ...
    protocols {
      pim-snooping{
        traceoptions {
          file pim.log size 10m;
          flag all;
          flag timer disable;
        }
      }
    }
  }
}
```

Example: Configuring PIM Snooping for VPLS explains the PIM snooping method. The use of the PIM proxying method is not discussed here and is outside the scope of this document. For more information about PIM proxying, see [PIM Snooping over VPLS](#).

Example: Configuring PIM Snooping for VPLS

IN THIS SECTION

- [Requirements | 1252](#)
- [Overview | 1252](#)
- [Configuration | 1253](#)
- [Verification | 1263](#)

This example shows how to configure PIM snooping in a virtual private LAN service (VPLS) to restrict multicast traffic to interested devices.

Requirements

This example uses the following hardware and software components:

- M Series Multiservice Edge Routers (M7i and M10i with Enhanced CFEB, M120, and M320 with E3 FPCs) or MX Series 5G Universal Routing Platforms (MX80, MX240, MX480, and MX960)
- Junos OS Release 13.2 or later

Overview

IN THIS SECTION

- [Topology | 1253](#)

The following example shows how to configure PIM snooping to restrict multicast traffic to interested devices in a VPLS.



NOTE: This example demonstrates PIM snooping by the use of a PIM snooping device to restrict multicast traffic. The use of the PIM proxying method to achieve PIM snooping is out of the scope of this document and is yet to be implemented in Junos OS.


Topology

In this example, two PE routers are connected to each other through a pseudowire connection. Router PE1 is connected to Routers CE1 and CE2. A multicast receiver is attached to Router CE2. Router PE2 is connected to Routers CE3 and CE4. A multicast source is connected to Router CE3, and a second multicast receiver is attached to Router CE4.

PIM snooping is configured on Routers PE1 and PE2. Hence, data sent from the multicast source is received only by members of the multicast group.

Figure 144 on page 1253 shows the topology used in this example.

Figure 144: PIM Snooping for VPLS



Configuration

IN THIS SECTION

- CLI Quick Configuration | 1254

- [Configuring PIM Snooping for VPLS | 1257](#)
- [Results | 1260](#)

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 PE1

```
set multicast-snooping-options traceoptions file snoop.log size 10m
set interfaces ge-2/0/0 encapsulation ethernet-vpls
set interfaces ge-2/0/0 unit 0 description toCE1
set interfaces ge-2/0/1 encapsulation ethernet-vpls
set interfaces ge-2/0/1 unit 0 description toCE2
set interfaces ge-2/0/2 unit 0 description toPE2
set interfaces ge-2/0/2 unit 0 family inet address 10.0.0.1/30
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.1/32
set routing-options router-id 10.255.1.1
set protocols mpls interface ge-2/0/1.0
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 10.255.1.1
set protocols bgp group toPE2 family l2vpn signaling
set protocols bgp group toPE2 neighbor 10.255.7.7
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-2/0/2.0
set protocols ldp interface lo0.0
set routing-instances titanium instance-type vpls
set routing-instances titanium vlan-id none
set routing-instances titanium interface ge-2/0/0.0
set routing-instances titanium interface ge-2/0/1.0
set routing-instances titanium route-distinguisher 101:101
set routing-instances titanium vrf-target target:201:201
set routing-instances titanium protocols vpls vpls-id 15
```



```
set routing-instances titanium protocols vpls site pe1 site-identifier 1
set routing-instances titanium protocols pim-snooping
```

Router CE1

```
set interfaces ge-2/0/0 unit 0 description toPE1
set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.10/30
set interfaces lo0 unit 0 family inet address 10.255.2.2./32
set routing-options router-id 10.255.2.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 10.255.3.3
set protocols pim interface all
```

Router CE2

```
set interfaces ge-2/0/0 unit 0 description toPE1
set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.6/30
set interfaces ge-2/0/1 unit 0 description toReceiver1
set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.13/30
set interfaces lo0 unit 0 family inet address 10.255.2.2
set routing-options router-id 10.255.2.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 10.255.3.3
set protocols pim interface all
```

Router PE2

```
set multicast-snooping-options traceoptions file snoop.log size 10m
set interfaces ge-2/0/0 encapsulation ethernet-vpls
set interfaces ge-2/0/0 unit 0 description toCE3
set interfaces ge-2/0/1 encapsulation ethernet-vpls
set interfaces ge-2/0/1 unit 0 description toCE4
set interfaces ge-2/0/2 unit 0 description toPE1
set interfaces ge-2/0/2 unit 0 family inet address 10.0.0.2/30
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.7.7/32
set routing-options router-id 10.255.7.7
set protocols mpls interface ge-2/0/2.0
set protocols bgp group toPE1 type internal
```

```

set protocols bgp group toPE1 local-address 10.255.7.7
set protocols bgp group toPE1 family l2vpn signaling
set protocols bgp group toPE1 neighbor 10.255.1.1
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp interface ge-2/0/2.0
set protocols ldp interface lo0.0
set routing-instances titanium instance-type vpls
set routing-instances titanium vlan-id none
set routing-instances titanium interface ge-2/0/0.0
set routing-instances titanium interface ge-2/0/1.0
set routing-instances titanium route-distinguisher 101:101
set routing-instances titanium vrf-target target:201:201
set routing-instances titanium protocols vpls vpls-id 15
set routing-instances titanium protocols vpls site pe2 site-identifier 2
set routing-instances titanium protocols pim-snooping

```

Router CE3 (RP)

```

set interfaces ge-2/0/0 unit 0 description toPE2
set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.18/30
set interfaces ge-2/0/1 unit 0 description toSource
set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.29/30
set interfaces lo0 unit 0 family inet address 10.255.3.3/32
set routing-options router-id 10.255.3.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp local address 10.255.3.3
set protocols pim interface all

```

Router CE4

```

set interfaces ge-2/0/0 unit 0 description toPE2
set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.22/30
set interfaces ge-2/0/1 unit 0 description toReceiver2
set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.25/30
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set routing-options router-id 10.255.4.4
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

```
set protocols pim rp static address 10.255.3.3
set protocols pim interface all
```

Configuring PIM Snooping for VPLS

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](#).



NOTE: This section includes a step-by-step configuration procedure for one or more routers in the topology. For comprehensive configurations for all routers, see "[CLI Quick Configuration](#)" on page 1254.

To configure PIM snooping for VPLS:

1. Configure the router interfaces forming the links between the routers.

Router PE2

```
[edit interfaces]
user@PE2# set ge-2/0/0 encapsulation ethernet-vpls
user@PE2# set ge-2/0/0 unit 0 description toCE3
user@PE2# set ge-2/0/1 encapsulation ethernet-vpls
user@PE2# set ge-2/0/1 unit 0 description toCE4
user@PE2# set ge-2/0/2 unit 0 description toPE1
user@PE2# set ge-2/0/2 unit 0 family mpls
user@PE2# set ge-2/0/2 unit 0 family inet address 10.0.0.2/30
user@PE2# set lo0 unit 0 family inet address 10.255.7.7/32
```



NOTE: ge-2/0/0.0 and ge-2/0/1.0 are configured as VPLS interfaces and connect to Routers CE3 and CE4. See *Virtual Private LAN Service User Guide* for more details.

Router CE3

```
[edit interfaces]
user@CE3# set ge-2/0/0 unit 0 description toPE2
user@CE3# set ge-2/0/0 unit 0 family inet address 10.0.0.18/30
user@CE3# set ge-2/0/1 unit 0 description toSource
```

```
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.0.29/30
user@CE3# set lo0 unit 0 family inet address 10.255.3.3/32
```



NOTE: The `ge-2/0/1.0` interface on Router CE3 connects to the multicast source.

Router CE4

```
[edit interfaces]
user@CE4# set ge-2/0/0 unit 0 description toPE2
user@CE4# set ge-2/0/0 unit 0 family inet address 10.0.0.22/30
user@CE4# set ge-2/0/1 unit 0 description toReceiver2
user@CE4# set ge-2/0/1 unit 0 family inet address 10.0.0.25/30
user@CE4# set lo0 unit 0 family inet address 10.255.4.4/32
```



NOTE: The `ge-2/0/1.0` interface on Router CE4 connects to a multicast receiver.

Similarly, configure Routers PE1, CE1, and CE2.

2. Configure the router IDs of all routers.

Router PE2

```
[edit routing-options]
user@PE2# set router-id 10.255.7.7
```

Similarly, configure other routers.

3. Configure an IGP on interfaces of all routers.

Router PE2

```
[edit protocols ospf area 0.0.0.0]
user@PE2# set interface ge-2/0/2.0
user@PE2# set interface lo0.0
```

Similarly, configure other routers.

4. Configure the LDP, MPLS, and BGP protocols on the PE routers.

```

Router PE2
[edit protocols]
user@PE2# set ldp interface lo0.0
user@PE2# set mpls interface ge-2/0/2.0
user@PE2# set bgp group toPE1 type internal
user@PE2# set bgp group toPE1 local-address 10.255.7.7
user@PE2# set bgp group toPE1 family l2vpn signaling
user@PE2# set bgp group toPE1 neighbor 10.255.1.1
user@PE2# set ldp interface ge-2/0/2.0

```

The BGP group is required for interfacing with the other PE router. Similarly, configure Router PE1.

5. Configure PIM on all CE routers.

Ensure that Router CE3 is configured as the rendezvous point (RP) and that the RP address is configured on other CE routers.

```

Router CE3
[edit protocols pim]
user@CE3# set rp local address 10.255.3.3
user@CE3# set interface all

```

```

Router CE4
[edit protocols pim]
user@CE4# set rp static address 10.255.3.3
user@CE4# set interface all

```

Similarly, configure Routers CE1 and CE2.

6. Configure multicast snooping options on the PE routers.

```

Router PE2
[edit multicast-snooping-options traceoptions]
user@PE2# set file snoop.log size 10m

```

Similarly, configure Router PE1.

7. Create a routing instance (titanium), and configure the VPLS on the PE routers.

```
Router PE2
[edit routing-instances titanium]
user@PE2# set instance-type vpls
user@PE2# set vlan-id none
user@PE2# set interface ge-2/0/0.0
user@PE2# set interface ge-2/0/1.0
user@PE2# set route-distinguisher 101:101
user@PE2# set vrf-target target:201:201
user@PE2# set protocols vpls vpls-id 15
user@PE2# set protocols vpls site pe2 site-identifier 2
```

Similarly, configure Router PE1.

8. Configure PIM snooping on the PE routers.

```
Router PE2
[edit routing-instances titanium]
user@PE2# set protocols pim-snooping
```

Similarly, configure Router PE1.

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, **show multicast-snooping-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-2/0/2 {
  unit 0 {
    description toPE1
    family inet {
      address 10.0.0.2/30;
    }
    family mpls;
  }
}
```

```
ge-2/0/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    description toCE3;
  }
}
ge-2/0/1 {
  encapsulation ethernet-vpls;
  unit 0 {
    description toCE4;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.7.7/32;
    }
  }
}
```

```
user@PE2# show routing-options
router-id 10.255.7.7;
```

```
user@PE2# show protocols
mpls {
  interface ge-2/0/2.0;
}
ospf {
  area 0.0.0.0 {
    interface ge-2/0/2.0;
    interface lo0.0;
  }
}
ldp {
  interface ge-2/0/2.0;
  interface lo0.0;
}
bgp {
  group toPE1 {
    type internal;
```

```

    local-address 10.255.7.7;
    family l2vpn {
        signaling;
    }
    neighbor 10.255.1.1;
}

```

```

user@PE2# show multicast-snooping-options
traceoptions {
    file snoop.log size 10m;
}

```

```

user@PE2# show routing-instances
titanium {
    instance-type vpls;
    vlan-id none;
    interface ge-2/0/0.0;
    interface ge-2/0/1.0;
    route-distinguisher 101:101;
    vrf-target target:201:201;
    protocols {
        vpls {
            site pe2 {
                site-identifier 2;
            }
            vpls-id 15;
        }
        pim-snooping;
    }
}

```

Similarly, confirm the configuration on all other routers. If you are done configuring the routers, enter `commit` from configuration mode.



NOTE: Use the `show protocols` command on the CE routers to verify the configuration for the PIM RP .

Verification

IN THIS SECTION

- [Verifying PIM Snooping for VPLS | 1263](#)

Confirm that the configuration is working properly.

Verifying PIM Snooping for VPLS

Purpose

Verify that PIM Snooping is operational in the network.

Action

To verify that PIM snooping is working as desired, use the following commands:

- **show pim snooping interfaces**
- **show pim snooping neighbors detail**
- **show pim snooping statistics**
- **show pim snooping join**
- **show pim snooping join extensive**
- **show multicast snooping route extensive instance *<instance-name>* group *<group-name>***

1. From operational mode on Router PE2, run the **show pim snooping interfaces** command.

```
user@PE2> show pim snooping interfaces
Instance: titanium

Learning-Domain: default

Name                State IP NbrCnt
ge-2/0/0.0          Up   4     1
ge-2/0/1.0          Up   4     1
```

```
DR address: 10.0.0.22
DR flooding is ON
```

The output verifies that PIM snooping is configured on the two interfaces connecting Router PE2 to Routers CE3 and CE4.

Similarly, check the PIM snooping interfaces on Router PE1.

2. From operational mode on Router PE2, run the **show pim snooping neighbors detail** command.

```
user@PE2> show pim snooping neighbors detail
Instance: titanium
Learning-Domain: default

Interface: ge-2/0/0.0

  Address: 10.0.0.18
    Uptime: 00:17:06
    Hello Option Holdtime: 105 seconds 99 remaining
    Hello Option DR Priority: 1
    Hello Option Generation ID: 552495559
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Tracking is supported

Interface: ge-2/0/1.0

  Address: 10.0.0.22
    Uptime: 00:15:16
    Hello Option Holdtime: 105 seconds 103 remaining
    Hello Option DR Priority: 1
    Hello Option Generation ID: 1131703485
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Tracking is supported
```

The output verifies that Router PE2 can detect the IP addresses of its PIM snooping neighbors (10.0.0.18 on CE3 and 10.0.0.22 on CE4).

Similarly, check the PIM snooping neighbors on Router PE1.

3. From operational mode on Router PE2, run the **show pim snooping statistics** command.

```
user@PE2> show pim snooping statistics
Instance: titanium
```

```

Learning-Domain: default

Tx J/P messages                0
RX J/P messages                246
Rx J/P messages -- seen        0
Rx J/P messages -- received    246
Rx Hello messages              1036
Rx Version Unknown             0
Rx Neighbor Unknown            0
Rx Upstream Neighbor Unknown   0
Rx J/P Busy Drop               0
Rx J/P Group Aggregate         0
Rx Malformed Packet            0

Rx No PIM Interface            0
Rx Bad Length                  0
Rx Unknown Hello Option        0
Rx Unknown Packet Type         0
Rx Bad TTL                     0
Rx Bad Destination Address     0
Rx Bad Checksum                 0
Rx Unknown Version             0

```

The output shows the number of hello and join/prune messages received by Router PE2. This verifies that PIM sparse mode is operational in the network.

4. Send multicast traffic from the source terminal attached to Router CE3, for the multicast group 203.0.113.1.
5. From operational mode on Router PE2, run the **show pim snooping join**, **show pim snooping join extensive**, and **show multicast snooping route extensive instance <instance-name> group <group-name>** commands to verify PIM snooping.

```

user@PE2> show pim snooping join
Instance: titanium
Learning-Domain: default

Group: 203.0.113.1
  Source: *
  Flags: sparse,rptree,wildcard
  Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0

```

```

Group: 203.0.113.1
  Source: 10.0.0.30
  Flags: sparse
  Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0

```

```

user@PE2> show pim snooping join extensive
Instance: titanium
Learning-Domain: default

Group: 203.0.113.1
  Source: *
  Flags: sparse,rptree,wildcard
  Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0
    Downstream port: ge-2/0/1.0
      Downstream neighbors:
        10.0.0.22 State: Join Flags: SRW Timeout: 180

Group: 203.0.113.1
  Source: 10.0.0.30
  Flags: sparse
  Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0
    Downstream port: ge-2/0/1.0
      Downstream neighbors:
        10.0.0.22 State: Join Flags: S Timeout: 180

```

The outputs show that multicast traffic sent for the group 203.0.113.1 is sent to Receiver 2 through Router CE4 and also display the upstream and downstream neighbor details.

```

user@PE2> show multicast snooping route extensive instance titanium group 203.0.113.1
Nexthop Bulking: OFF

      Family: INET

Group: 203.0.113.1/24
  Bridge-domain: titanium
  Mesh-group: __all_ces__
  Downstream interface list:
    ge-2/0/1.0 -(1072)
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 1048577
  Route state: Active

```

```

Forwarding state: Forwarding

Group: 203.0.113.1/24
Source: 10.0.0.8
Bridge-domain: titanium
Mesh-group: __all_ces__
Downstream interface list:
    ge-2/0/1.0 -(1072)
Statistics: 0 kBps, 0 pps, 0 packets
Next-hop ID: 1048577
Route state: Active
Forwarding state: Forwarding

```

Meaning

PIM snooping is operational in the network.


IGMP and MLD Snooping for VPLS

IN THIS SECTION

- [Handling of L3 Routes with Integrated Routing and Bridging \(IRB\) within VPLS | 1271](#)

You can enable IGMP or MLD snooping in a virtual private LAN service (VPLS) to ensure that the customer-facing interfaces receive only the multicast traffic it has requested for. This snooping can be enabled with or without Integrated routing and bridging (IRB).

Figure 145: Basic VPLS topology



A logical full mesh of all participating Provider Edge (PE) routers is necessary for IGMP/MLD snooping to work in VPLS. In other words, every PE router is connected to every other PE router by a pseudowire resulting in a full mesh infrastructure. When you enable IGMP/MLD snooping over VPLS, multicast traffic is forwarded to all pseudowires that receive IGMP/MLD reports from remote (PE) devices. IGMP/MLD membership queries and join reports are flooded to all the pseudowires belonging to that VPLS instance. This allows for optimization of the multicast data flow to only those members of the group that are interested. The operating system builds a database of group members per service by listening to IGMP/MLD queries and reports from each PE device.



NOTE:

- VPLS multicast traffic forwarded from the core to access is based on the routes learnt via IGMP or MLD snooping.
- VPLS multicast traffic from access is flooded to the core even when there are no remote receivers.

IGMPv2/v3 snooping is supported in VPLS for IPv4 multicast traffic. To configure IGMP snooping on a PE router, include the `igmp-snooping` statement at the [edit routing-instances *instance-name* protocols] hierarchy level:

```
routing-instances {
  vpls1 {
    instance-type virtual-switch;
    protocols {
      igmp-snooping {
        vlan <vlan_name>
          traceoptions {
            file ...;
            flag [all | route | normal | general | state | policy | task | timer | packets
| query | report | leave]
              [detail | disable | receive | send];
          }
          l2-querier {
            source-address <ip-address>;
          }
          proxy {
            source-address <ip-address>;
          }
          query-interval <seconds>;
          query-last-member-interval <1..1024 seconds>;
          query-response-interval <seconds>;
        }
      }
    }
  }
}
```

```

robust-count <2..10>;
immediate-leave;
interface <interface-name> {
    multicast-router-interface;
    host-only-interface;
    group-limit <limit>;
    static {
        group <ip-address>;
        group <ip-address> {
            source <ip-address>;
        }
    }
}
}
}
}
}
}
}
}
}
}

```

Similarly, MLDv1/v2 snooping is supported in VPLS for IPv6 multicast traffic. To configure MLD snooping on a PE router, include the `mld-snooping` statement at the [edit routing-instances *instance-name* protocols] hierarchy level:

```

routing-instances {
  vpls1 {
    instance-type virtual-switch;
    protocols {
      mld-snooping {
        vlan <vlan_name>
        traceoptions {
          file ...;
          flag [all | client-notification | general | group | host-notification | leave |
normal | packets | policy | query | report | route | state | task | timer]
          [detail | disable | receive | send];
        }
        immediate-leave;
        query-interval <seconds>;
        query-last-member-interval <seconds>;
        query-response-interval <seconds>;
        robust-count <count>;
        proxy {
          source-address <ipv6-address>;
        }
      }
    }
  }
}

```

```

interface <interface-name> {
    host-only-interface;
    immediate-leave;
    multicast-router-interface;
    group-limit <max-number-of-groups>;
    static {
        group <ipv6-address> {
            source <ipv6-address>;
        }
    }
}
}
}
}
}
}
}

```



NOTE: MLDv2 requires specific hardware database profiles to allocate tables with different sizes in the hardware. To configure MLD v2 within a VPLS instance, include the `balanced-exem` option or the `l3-xl` option at the [edit system packet-forwarding-options hw-db-profile] hierarchy level.

To configure IRB within a VPLS instance, include the `l3-interface` *irb-interface-name* statement at the [edit routing-instances *routing-instance-name* instance-type virtual-switch] hierarchy level:

```

routing-instances {
    vpls1 {
        instance-type virtual-switch;
        vlans {
            vlan<id> {
                l3-interface irb.0;
            }
        }
    }
}
}

```

If the `no-local-switching` statement is configured under the [edit bridge-domains *bridge-domain-name*] hierarchy level, frames arriving on a CE interface are sent to a VPLS edge (VE) device or core-facing interfaces only. This ensures that the access ports in the VPLS domain do not forward packets to each other.

To configure the UNI logical interface, the `vlan-bridge` option must be included under the `[edit interfaces interface-name unit logical-unit-number encapsulation]` hierarchy level:

```
et-0/0/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 400;
  }
}
```




Configuration of VPLS ports is supported using the `virtual-switch` routing instance. Routing instance of type `vpls` is not supported.

```
routing-instances {
  vpls1 {
    instance-type virtual-switch;
    protocols {
      vpls {
        neighbor 10.255.67.22;
        no-tunnel-services;
        vpls-id 200;
      }
    }
  }
}
```

Handling of L3 Routes with Integrated Routing and Bridging (IRB) within VPLS

[Figure 146 on page 1272](#) illustrates possible L3 routing cases in a PE with PIM and IRB enabled.

Figure 146: L3 routes in VPLS



NOTE: Disable routing on all PE routers except for the centralized PE router to avoid traffic loops.

- The source is external (L3 interface) and the receiver is in a VPLS domain. The IGMP/MLD queries are reinjected into the multicast router and LSI port. The multicast traffic is encapsulated and routed via the pseudowires.
- The source is in a VPLS domain and the receiver is external (L3 interface). The VPLS multicast traffic is decapsulated and routed to the external interface.
- The source is in a VPLS domain and the receiver is in the same VPLS. The VPLS multicast traffic is bridged within the VPLS bridge domain.
- The source is in a VPLS domain and the receiver is in a different VPLS domain. The VPLS multicast traffic is routed across VPLS circuits and IRBs.

Configure Multicast Routing Options

IN THIS CHAPTER

- [Examples: Configuring Administrative Scoping | 1273](#)
- [Examples: Configuring Bandwidth Management | 1284](#)
- [Examples: Configuring the Multicast Forwarding Cache | 1317](#)
- [Example: Configuring Ingress PE Redundancy | 1327](#)

Examples: Configuring Administrative Scoping

IN THIS SECTION

- [Understanding Multicast Administrative Scoping | 1273](#)
- [Example: Creating a Named Scope for Multicast Scoping | 1275](#)
- [Example: Using a Scope Policy for Multicast Scoping | 1279](#)
- [Example: Configuring Externally Facing PIM Border Routers | 1283](#)

Understanding Multicast Administrative Scoping

You use multicast scoping to limit multicast traffic by configuring it to an administratively defined topological region. Multicast scoping controls the propagation of multicast messages—both multicast group join messages that are sent upstream toward a source and data forwarding downstream. Scoping can relieve stress on scarce resources, such as bandwidth, and improve privacy or scaling properties.

IP multicast implementations can achieve some level of scoping by using the time-to-live (TTL) field in the IP header. However, TTL scoping has proven difficult to implement reliably, and the resulting schemes often are complex and difficult to understand.

Administratively scoped IP multicast provides clearer and simpler semantics for multicast scoping. Packets addressed to administratively scoped multicast addresses do not cross configured administrative

boundaries. Administratively scoped multicast addresses are locally assigned, and hence are not required to be unique across administrative boundaries.

The administratively scoped IP version 4 (IPv4) multicast address space is the range from 239.0.0.0 through 239.255.255.255.

The structure of the IPv4 administratively scoped multicast space is based loosely on the IP version 6 (IPv6) addressing architecture described in RFC 1884, *IP Version 6 Addressing Architecture*.

There are two well-known scopes:

- IPv4 local scope—This scope comprises addresses in the range 239.255.0.0/16. The local scope is the minimal enclosing scope and is not further divisible. Although the exact extent of a local scope is site-dependent, locally scoped regions must not span any other scope boundary and must be contained completely within or be equal to any larger scope. If scope regions overlap in an area, the area of overlap must be within the local scope.
- IPv4 organization local scope—This scope comprises 239.192.0.0/14. It is the space from which an organization allocates subranges when defining scopes for private use.

The ranges 239.0.0.0/10, 239.64.0.0/10, and 239.128.0.0/10 are unassigned and available for expansion of this space.

Two other scope classes already exist in IPv4 multicast space: the statically assigned link-local scope, which is 224.0.0.0/24, and the static global scope allocations, which contain various addresses.

All scoping is inherently bidirectional in the sense that join messages and data forwarding are controlled in both directions on the scoped interface.

You can configure multicast scoping either by creating a named scope associated with a set of routing device interfaces and an address range, or by referencing a scope policy that specifies the interfaces and configures the address range as a series of filters. You cannot combine the two methods (the commit operation fails for a configuration that includes both). The methods differ somewhat in their requirements and result in different output from the `show multicast scope` command.

Routing loops must be avoided in IP multicast networks. Because multicast routers must replicate packets for each downstream branch, not only do looping packets not arrive at a destination, but each pass around the loop multiplies the number of looping packets, eventually overwhelming the network.

Scoping limits the routers and interfaces that can be used to forward a multicast packet. Scoping can use the TTL field in the IP packet header, but TTL scoping depends on the administrator having a thorough knowledge of the network topology. This topology can change as links fail and are restored, making TTL scoping a poor solution for multicast.

Multicast scoping is administrative in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365. Routers at the boundary must be able to filter multicast packets and make sure that the packets do not stray beyond the established limit.

Administrative scoping is much better than TTL scoping, but in many cases the dropping of administratively scoped packets is still determined by the network administrator. For example, the multicast address range 239/8 is defined in RFC 2365 as administratively scoped, and packets using this range are not to be forwarded beyond a network “boundary,” usually a routing domain. But only the network administrator knows where the border routers are and can implement the scoping correctly.

Multicast groups used by unicast routing protocols, such as 224.0.0.5 for all OSPF routers, are administratively scoped for that LAN only. This scoping allows the same multicast address to be used without conflict on every LAN running OSPF.

SEE ALSO

[Supported IP Multicast Protocol Standards | 21](#)

[Standards Reference](#)

Example: Creating a Named Scope for Multicast Scoping

IN THIS SECTION

- [Requirements | 1275](#)
- [Overview | 1276](#)
- [Configuration | 1276](#)
- [Verification | 1279](#)

This example shows how to configure multicast scoping with four scopes: **local**, **organization**, **engineering**, and **marketing**.

Requirements

Before you begin:

- Configure a tunnel interface. See the [Junos OS Network Interfaces Library for Routing Devices](#).
- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview

The **local** scope is configured on a GRE tunnel interface. The **organization** scope is configured on a GRE tunnel interface and a SONET/SDH interface. The **engineering** scope is configured on an IP-IP tunnel interface and two SONET/SDH interfaces. The **marketing** scope is configured on a GRE tunnel interface and two SONET/SDH interfaces. The Junos OS can scope any user-configurable IPv6 or IPv4 group.

To configure multicast scoping by defining a named scope, you must specify a name for the scope, the set of routing device interfaces on which you are configuring scoping, and the scope's address range.



NOTE: The prefix specified with the `prefix` statement must be unique for each scope statement. If multiple scopes contain the same prefix, only the last scope applies to the interfaces. If you need to scope the same prefix on multiple interfaces, list all of them in the interface statement for a single scope statement.

When you configure multicast scoping with a named scope, all scope boundaries must include the **local** scope. If this scope is not configured, it is added automatically at all scoped interfaces. The **local** scope limits the use of the multicast group **239.255.0.0/16** to an attached LAN.

Configuration

IN THIS SECTION

- Procedure | [1276](#)
- Results | [1278](#)

Procedure

CLI Quick Configuration

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

```
set routing-options multicast scope local prefix fe00::239.255.0.0/128
set routing-options multicast scope local interface gr-2/1/0.0
set routing-options multicast scope organization prefix 239.192.0.0/14
set routing-options multicast scope organization interface gr-2/1/0.0
```

```

set routing-options multicast scope organization interface so-0/0/0.0
set routing-options multicast scope engineering prefix 239.255.255.0/24
set routing-options multicast scope engineering interface ip-2/1/0.0
set routing-options multicast scope engineering interface so-0/0/1.0
set routing-options multicast scope engineering interface so-0/0/2.0
set routing-options multicast scope marketing prefix 239.255.254.0/24
set routing-options multicast scope marketing interface gr-2/1/0.0
set routing-options multicast scope marketing interface so-0/0/2.0
set routing-options multicast scope marketing interface so-1/0/0.0

```

Step-by-Step Procedure

1. 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 [Junos OS CLI User Guide](#).

Configure the local scope.

```

[edit routing-options multicast]
user@host# set scope local interface gr-2/1/0
user@host# set scope localprefix fe00::239.255.0.0/128

```

2. Configure the organization scope.

```

[edit routing-options multicast]
user@host# set scope organization interface [ gr-2/1/0 so-0/0/0 ]
user@host# set scope organization prefix 239.192.0.0/14

```

3. Configure the engineering scope.

```

[edit routing-options multicast]
user@host# set scope engineering interface [ ip-2/1/0 so-0/0/1 so-0/0/2 ]
user@host# set scope engineering prefix 239.255.255.0/24

```

4. Configure the marketing scope.

```
[edit routing-options multicast]
user@host# set scope marketing interface [ gr-2/1/0 so-0/0/2 so-1/0/0 ]
user@host# set scope marketing prefix 239.255.254.0/24
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show routing-options` command.

```
user@host# show routing-options
multicast {
  scope local {
    interface gr-2/1/0;
    prefix fe00::239.255.0.0/128;
  }
  scope organization {
    interface [ gr-2/1/0 so-0/0/0 ];
    prefix 239.192.0.0/14;
  }
  scope engineering {
    interface [ ip-2/1/0 so-0/0/1 so-0/0/2 ];
    prefix 239.255.255.0/24;
  }
  scope marketing {
    interface [ gr-2/1/0 so-0/0/2 so-1/0/0 ];
    prefix 239.255.254.0/24;
  }
}
```


Verification

To verify that group scoping is in effect, issue the `show multicast scope` command:

```
user@host> show multicast scope
Resolve
Scope name      Group prefix      Interface      Rejects
local           fe00::239.255.0.0/128 gr-2/1/00
organization    239.192.0.0/14    gr-2/1/0      so-0/0/00
engineering     239.255.255.0/24 ip-2/1/0       so-0/0/1 so-0/0/20
marketing       239.255.254.0/24 gr-2/1/0       so-0/0/2 so-1/0/00
```

When you configure scoping with a named scope, the `show multicast scope` operational mode command displays the names of the defined scopes, prefixes, and interfaces.

Example: Using a Scope Policy for Multicast Scoping

IN THIS SECTION

- [Requirements | 1279](#)
- [Overview | 1279](#)
- [Configuration | 1280](#)
- [Verification | 1283](#)

This example shows how to configure a multicast scope policy named **allow-auto-rp-on-backbone**, allowing packets for auto-RP groups 224.0.1.39/32 and 224.0.1.40/32 on backbone-facing interfaces, and rejecting all other addresses in the 224.0.1.0/24 and 239.0.0.0/8 address ranges.

Requirements

Before you begin:

- Configure an interior gateway protocol or static routing. See the [Junos OS Routing Protocols Library for Routing Devices](#).

Overview

Each referenced policy must be correctly configured at the `[edit policy-options]` hierarchy level, specifying the set of routing device interfaces on which to configure scoping, and defining the scope's

address range as a series of route filters. Only the **interface**, **route-filter**, and **prefix-list** match conditions are supported for multicast scope policies. All other configured match conditions are ignored. The only actions supported are **accept**, **reject**, and the policy flow actions **next-term** and **next-policy**. The **reject** action means that joins and multicast forwarding are suppressed in both directions on the configured interfaces. The **accept** action allows joins and multicast forwarding in both directions on the interface. By default, scope policies apply to all interfaces. The default action is **accept**.



NOTE: Multicast scoping configured with a scope policy differs in some ways from scoping configured with a named scope (which uses the `scope` statement):

- You cannot apply a scope policy to a specific routing instance, because all scope policies apply to all routing instances. In contrast, a named scope does apply individually to a specific routing instance.
- In contrast to scoping with a named scope, scoping with a scope policy does not automatically add the **local** scope at scope boundaries. You must explicitly configure the local scope boundaries. The **local** scope limits the use of the multicast group 239.255.0.0/16 to an attached LAN.

Configuration

IN THIS SECTION

- [Procedure | 1280](#)
- [Results | 1282](#)

Procedure

CLI Quick Configuration

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

```
set policy-options policy-statement allow-auto-rp-on-backbone term allow-auto-rp from interface so-0/0/0.0
set policy-options policy-statement allow-auto-rp-on-backbone term allow-auto-rp from interface so-0/0/1.0
```

```

set policy-options policy-statement allow-auto-rp-on-backbone term allow-auto-rp from route-
filter 224.0.1.39/32 exact
set policy-options policy-statement allow-auto-rp-on-backbone term allow-auto-rp from route-
filter 224.0.1.40/32 exact
set policy-options policy-statement allow-auto-rp-on-backbone term allow-auto-rp then accept
set policy-options policy-statement allow-auto-rp-on-backbone term reject-these from route-
filter 224.0.1.0/24 orlonger
set policy-options policy-statement allow-auto-rp-on-backbone term reject-these from route-
filter 239.0.0.0/8 orlonger
set policy-options policy-statement allow-auto-rp-on-backbone term reject-these then reject
set routing-options multicast scope-policy allow-auto-rp-on-backbone

```

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 [Junos OS CLI User Guide](#).

1. Define which packets are allowed.

```

[edit policy-options policy-statement allow-auto-rp-on-backbone]
user@host# set term allow-auto-rp from interface so-0/0/0.0
user@host# set term allow-auto-rp from interface so-0/0/1.0
user@host# set term allow-auto-rp from route-filter 224.0.1.39/32 exact
user@host# set term allow-auto-rp from route-filter 224.0.1.40/32 exact
user@host# set term allow-auto-rp then accept

```

2. Define which packets are not allowed.

```

[edit policy-options policy-statement allow-auto-rp-on-backbone]
user@host# set term reject-these from route-filter 224.0.1.0/24 orlonger
user@host# set term reject-these from route-filter 239.0.0.0/8 orlonger
user@host# set term reject-these then reject

```

3. Apply the policy.

```

[edit routing-options multicast]
user@host# set scope-policy allow-auto-rp-on-backbone

```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the `show policy-options` and `show routing-options` commands.

```
user@host# show policy-options
policy-statement allow-auto-rp-on-backbone {
  term allow-auto-rp {
    from {
      /* backbone-facing interfaces */
      interface [ so-0/0/0.0 so-0/0/1.0 ];
      route-filter 224.0.1.39/32 exact;
      route-filter 224.0.1.40/32 exact;
    }
    then {
      accept;
    }
  }
  term reject-these {
    from {
      route-filter 224.0.1.0/24 orlonger;
      route-filter 239.0.0.0/8 orlonger;
    }
    then reject;
  }
}
```

```
user@host# show routing-options
multicast {
  scope-policy allow-auto-rp-on-backbone;
}
```

Verification

To verify that the scope policy is in effect, issue the `show multicast scope configuration mode` command:

```
user@host> show multicast scope
Scope policy: [ allow-auto-rp-on-backbone ]
```

When you configure multicast scoping with a scope policy, the `show multicast scope operational mode` command displays only the name of the scope policy.

Example: Configuring Externally Facing PIM Border Routers

In this example, you add the `scope` statement at the `[edit routing-options multicast]` hierarchy level to prevent auto-RP traffic from “leaking” into or out of your PIM domain. Two scopes defined below, `auto-rp-39` and `auto-rp-40`, are for specific addresses. The `scoped-range` statement defines a group range, thus preventing group traffic from leaking.

```
routing-options {
  multicast {
    scope auto-rp-39 {
      prefix 224.0.1.39/32;
      interface t1-0/0/0.0;
    }
    scope auto-rp-40 {
      prefix 224.0.1.40/32;
      interface t1-0/0/0.0;
    }
    scope scoped-range {
      prefix 239.0.0.0/8;
      interface t1-0/0/0.0;
    }
  }
}
```

RELATED DOCUMENTATION

[Examples: Configuring Bandwidth Management | 1284](#)

[Examples: Configuring the Multicast Forwarding Cache | 1317](#)

Examples: Configuring Bandwidth Management

IN THIS SECTION

- [Understanding Bandwidth Management for Multicast | 1284](#)
- [Bandwidth Management and PIM Graceful Restart | 1285](#)
- [Bandwidth Management and Source Redundancy | 1285](#)
- [Logical Systems and Bandwidth Oversubscription | 1285](#)
- [Example: Defining Interface Bandwidth Maximums | 1287](#)
- [Example: Configuring Multicast with Subscriber VLANs | 1291](#)
- [Configuring Multicast Routing over IP Demux Interfaces | 1308](#)
- [Classifying Packets by Egress Interface | 1309](#)
- [Recycle Bandwidth Management | 1312](#)
- [Example: Configuring Recycle Bandwidth | 1314](#)

Understanding Bandwidth Management for Multicast

Bandwidth management enables you to control the multicast flows that leave a multicast interface. This control enables you to better manage your multicast traffic and reduce or eliminate the chances of interface oversubscription or congestion.

Bandwidth management ensures that multicast traffic oversubscription does not occur on an interface. When managing multicast bandwidth, you define the maximum amount of multicast bandwidth that an individual interface can use as well as the bandwidth individual multicast flows use.

For example, the routing software cannot add a flow to an interface if doing so exceeds the allowed bandwidth for that interface. Under these circumstances, the interface is rejected. This rejection, however, does not prevent a multicast protocol (for example, PIM) from sending a join message upstream. Traffic continues to arrive on the router, even though the router is not sending the flow from the expected outgoing interfaces.

You can configure the flow bandwidth statically by specifying a bandwidth value for the flow in bits per second, or you can enable the flow bandwidth to be measured and adaptively changed. When using the adaptive bandwidth option, the routing software queries the statistics for the flows to be measured at 5-second intervals and calculates the bandwidth based on the queries. The routing software uses the maximum value measured within the last minute (that is, the last 12 measuring points) as the flow bandwidth.

For more information, see the following sections:

- ["Bandwidth Management and PIM Graceful Restart" on page 1285](#)
- ["Bandwidth Management and Source Redundancy" on page 1285](#)
- ["Logical Systems and Bandwidth Oversubscription" on page 1285](#)

Bandwidth Management and PIM Graceful Restart

When using PIM graceful restart, after the routing process restarts on the Routing Engine, previously admitted interfaces are always readmitted and the available bandwidth is adjusted on the interfaces. When using the adaptive bandwidth option, the bandwidth measurement is initially based on the configured or default starting bandwidth, which might be inaccurate during the first minute. This means that new flows might be incorrectly rejected or admitted temporarily. You can correct this problem by issuing the **clear multicast bandwidth-admission** operational command.

If PIM graceful restart is not configured, after the routing process restarts, previously admitted or rejected interfaces might be rejected or admitted in an unpredictable manner.

SEE ALSO

| [CLI Explorer](#)

Bandwidth Management and Source Redundancy

When using source redundancy, multiple sources (for example, s1 and s2) might exist for the same destination group (g). However, only one of the sources can actively transmit at any time. In this case, multiple forwarding entries—(s1,g) and (s2,g)—are created after each goes through the admission process.

With redundant sources, unlike unrelated entries, an OIF that is already admitted for one entry—for example, (s1,g)—is automatically admitted for other redundancy entries—for example, (s2,g). The remaining bandwidth on the interface is deducted each time an outbound interface is added, even though only one sender actively transmits. By measuring bandwidth, the bandwidth deducted for the inactive entries is credited back when the router detects no traffic is being transmitted.

For more information about defining redundant sources, see ["Example: Configuring a Multicast Flow Map" on page 1321](#).

Logical Systems and Bandwidth Oversubscription

You can manage bandwidth at both the physical and *logical interface* level. However, if more than one logical system shares the same physical interface, the interface might become oversubscribed.

Oversubscription occurs if the total bandwidth of all separately configured maximum bandwidth values for the interfaces on each logical system exceeds the bandwidth of the physical interface.

When displaying interface bandwidth information, a negative available bandwidth value indicates oversubscription on the interface.

Interface bandwidth can become oversubscribed when the configured maximum bandwidth decreases or when some flow bandwidths increase because of a configuration change or an actual increase in the traffic rate.

Interface bandwidth can become available again if one of the following occurs:

- The configured maximum bandwidth increases.
- Some flows are no longer transmitted from interfaces, and bandwidth reserves for them are now available to other flows.
- Some flow bandwidths decrease because of a configuration change or an actual decrease in the traffic rate.

Interfaces that are rejected for a flow because of insufficient bandwidth are not automatically readmitted, even when bandwidth becomes available again. Rejected interfaces have an opportunity to be readmitted when one of the following occurs:

- The multicast routing protocol updates the forwarding entry for the flow after receiving a join, leave, or prune message or after a topology change occurs.
- The multicast routing protocol updates the forwarding entry for the flow due to configuration changes.
- You manually reapply bandwidth management to a specific flow or to all flows using the **clear multicast bandwidth-admission** operational command.

In addition, even if previously available bandwidth is no longer available, already admitted interfaces are not removed until one of the following occurs:

- The multicast routing protocol explicitly removes the interfaces after receiving a leave or prune message or after a topology change occurs.
- You manually reapply bandwidth management to a specific flow or to all flows using the **clear multicast bandwidth-admission** operational command.

SEE ALSO

| [CLI Explorer](#)

Example: Defining Interface Bandwidth Maximums

IN THIS SECTION

- Requirements | 1287
- Overview | 1287
- Configuration | 1289
- Verification | 1291

This example shows you how to configure the maximum bandwidth for a physical or logical interface.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure a multicast protocol. This feature works with the following multicast protocols:
 - DVMRP
 - PIM-DM
 - PIM-SM
 - PIM-SSM

Overview

IN THIS SECTION

- Topology | 1288

The maximum bandwidth setting applies admission control either against the configured interface bandwidth or against the native speed of the underlying interface (when there is no configured bandwidth for the interface).

If you configure several logical interfaces (for example, to support VLANs or PVCs) on the same underlying physical interface, and no bandwidth is configured for the logical interfaces, it is assumed that the logical interfaces all have the same bandwidth as the underlying interface. This can cause oversubscription. To prevent oversubscription, configure bandwidth for the logical interfaces, or configure admission control at the physical interface level.

You only need to define the maximum bandwidth for an interface on which you want to apply bandwidth management. An interface that does not have a defined maximum bandwidth transmits all multicast flows as determined by the multicast protocol that is running on the interface (for example, PIM).

If you specify **maximum-bandwidth** without including a bits-per-second value, admission control is enabled based on the bandwidth configured for the interface. In the following example, admission control is enabled for logical interface unit **200**, and the maximum bandwidth is 20 Mbps. If the bandwidth is not configured on the interface, the maximum bandwidth is the link speed.

```
routing-options {
  multicast {
    interface fe-0/2/0.200 {
      maximum-bandwidth;
    }
  }
  interfaces {
    fe-0/2/0 {
      unit 200 {
        bandwidth 20m;
      }
    }
  }
}
```

Topology

Configuration

IN THIS SECTION

- Procedure | 1289
- Results | 1290

Procedure

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.

```
set interfaces fe-0/2/0 unit 200 bandwidth 20m
set routing-options multicast interface fe-0/2/0.200 maximum-bandwidth
set routing-options multicast interface fe-0/2/1 maximum-bandwidth 60m
set routing-options multicast interface fe-0/2/1.200 maximum-bandwidth 10m
```

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 [Junos OS CLI User Guide](#).

To configure a bandwidth maximum:

1. Configure the a logical interface bandwidth.

```
[edit interfaces]
user@host# set fe-0/2/0 unit 200 bandwidth 20m
```

2. Enable admission control on the logical interface.

```
[edit routing-options]
user@host# set multicast interface fe-0/2/0.200 maximum-bandwidth
```

3. On a physical interface, enable admission control and set the maximum bandwidth to 60 Mbps.

```
[edit routing-options]
user@host# set multicast interface fe-0/2/1 maximum-bandwidth 60m
```

4. For a logical interface on the same physical interface shown in Step "3" on page 1290, set a smaller maximum bandwidth.

```
[edit routing-options]
user@host# set multicast interface fe-0/2/1.200 maximum-bandwidth 10m
```

Results

Confirm your configuration by entering the **show interfaces** and **show routing-options** commands.

```
user@host# show interfaces
fe-0/2/0 {
  unit 200 {
    bandwidth 20m;
  }
}
```

```
user@host# show routing-options
multicast {
  interface fe-0/2/0.200 {
    maximum-bandwidth;
  }
  interface fe-0/2/1 {
    maximum-bandwidth 60m;
  }
  interface fe-0/2/1.200 {
    maximum-bandwidth 10m;
  }
}
```

```
}  
}
```

Verification

To verify the configuration, run the **show multicast interface** command.

SEE ALSO

[Example: Configuring a Multicast Flow Map | 1321](#)

Example: Configuring Multicast with Subscriber VLANs

IN THIS SECTION

- [Requirements | 1291](#)
- [Overview and Topology | 1292](#)
- [Configuration | 1296](#)
- [Verification | 1307](#)

This example shows how to configure an MX Series router to function as a broadband service router (BSR).

Requirements

This example uses the following hardware components:

- One MX Series router or EX Series switch with a PIC that supports traffic control profile queuing
- One DSLAM

Before you begin:

- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure PIM and IGMP or MLD on the interfaces.

Overview and Topology

IN THIS SECTION

- [Topology | 1295](#)

When multiple BSR interfaces receive IGMP and MLD join and leave requests for the same multicast stream, the BSR sends a copy of the multicast stream on each interface. Both the multicast control packets (IGMP and MLD) and the multicast data packets flow on the same BSR interface, along with the unicast data. Because all per-customer traffic has its own interface on the BSR, per-customer accounting, call admission control (CAC), and quality-of-service (QoS) adjustment are supported. The QoS bandwidth used by multicast reduces the unicast bandwidth.

Multiple interfaces on the BSR might connect to a shared device (for example, a DSLAM). The BSR sends the same multicast stream multiple times to the shared device, thus wasting bandwidth. It is more efficient to send the multicast stream one time to the DSLAM and replicate the multicast streams in the DSLAM. There are two approaches that you can use.

The first approach is to continue to send unicast data on the per-customer interfaces, but have the DSLAM route all the per-customer IGMP and MLD join and leave requests to the BSR on a single dedicated interface (a multicast VLAN). The DSLAM receives the multicast streams from the BSR on the dedicated interface with no unnecessary replication and performs the necessary replication to the customers. Because all multicast control and data packets use only one interface, only one copy of a stream is sent even if there are multiple requests. This approach is called reverse outgoing interface (OIF) mapping. Reverse OIF mapping enables the BSR to propagate the multicast state of the shared interface to the customer interfaces, which enables per-customer accounting and QoS adjustment to work. When a customer changes the TV channel, the router gateway (RG) sends an IGMP or MLD join and leave messages to the DSLAM. The DSLAM transparently passes the request to the BSR through the multicast VLAN. The BSR maps the IGMP or MLD request to one of the subscriber VLANs based on the IP source address or the source MAC address. When the subscriber VLAN is found, QoS adjustment and accounting are performed on that VLAN or interface.

The second approach is for the DSLAM to continue to send unicast data and all the per-customer IGMP and MLD join and leave requests to the BSR on the individual customer interfaces, but to have the multicast streams arrive on a single dedicated interface. If multiple customers request the same multicast stream, the BSR sends one copy of the data on the dedicated interface. The DSLAM receives the multicast streams from the BSR on the dedicated interface and performs the necessary replication to the customers. Because the multicast control packets use many customer interfaces, configuration on the BSR must specify how to map each customer's multicast data packets to the single dedicated output interface. QoS adjustment is supported on the customer interfaces. CAC is supported on the shared interface. This second approach is called multicast OIF mapping.

OIF mapping and reverse OIF mapping are not supported on the same customer interface or shared interface. This example shows how to configure the two different approaches. Both approaches support QoS adjustment, and both approaches support MLD/IPv6. The reverse OIF mapping example focuses on IGMP/IPv4 and enables QoS adjustment. The OIF mapping example focuses on MLD/IPv6 and disables QoS adjustment.

The first approach (reverse OIF mapping) includes the following statements:

- **flow-map**—Defines a flow map that controls the bandwidth for each flow.
- **maximum-bandwidth**—Enables CAC.
- **reverse-oif-mapping**—Enables the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD join or leave request that it receives over the multicast VLAN.

After the subscriber VLAN is identified, the routing device immediately adjusts the QoS (in this case, the bandwidth) on that VLAN based on the addition or removal of a subscriber.

The routing device uses IGMP and MLD join or leave reports to obtain the subscriber VLAN information. This means that the connecting equipment (for example, the DSLAM) must forward all IGMP and MLD reports to the routing device for this feature to function properly. Using report suppression or an IGMP proxy can result in reverse OIF mapping not working properly.

- **subscriber-leave-timer**—Introduces a delay to the QoS update. After receiving an IGMP or MLD leave request, this statement defines a time delay (between 1 and 30 seconds) that the routing device waits before updating the QoS for the remaining subscriber interfaces. You might use this delay to decrease how often the routing device adjusts the overall QoS bandwidth on the VLAN when a subscriber sends rapid leave and join messages (for example, when changing channels in an IPTV network).
- **traffic-control-profile**—Configures a shaping rate on the logical interface. The configured shaping rate must be configured as an absolute value, not as a percentage.

The second approach (OIF mapping) includes the following statements:

- **map-to-interface**—In a policy statement, enables you to build the OIF map.

The OIF map is a routing policy statement that can contain multiple terms. When creating OIF maps, keep the following in mind:

- If you specify a physical interface (for example, **ge-0/0/0**), a ".0" is appended to the interface to create a logical interface (for example, **ge-0/0/0.0**).
- Configure a routing policy for each logical system. You cannot configure routing policies dynamically.
- The interface must also have IGMP, MLD, or PIM configured.

- You cannot map to a mapped interface.
- We recommend that you configure policy statements for IGMP and MLD separately.
- Specify either a logical interface or the keyword **self**. The **self** keyword specifies that multicast data packets be sent on the same interface as the control packets and that no mapping occur. If no term matches, then no multicast data packets are sent.
- **no-qos-adjust**—Disables QoS adjustment.

QoS adjustment decreases the available bandwidth on the client interface by the amount of bandwidth consumed by the multicast streams that are mapped from the client interface to the shared interface. This action always occurs unless it is explicitly disabled.

If you disable QoS adjustment, available bandwidth is not reduced on the customer interface when multicast streams are added to the shared interface.



NOTE: You can dynamically disable QoS adjustment for IGMP and MLD interfaces using dynamic profiles.

- **oif-map**—Associate a map with an IGMP or MLD interface. The OIF map is then applied to all IGMP or MLD requests received on the configured interface. In this example, subscriber VLANs 1 and 2 have MLD configured, and each VLAN points to an OIF map that directs some traffic to **ge-2/3/9.4000**, some traffic to **ge-2/3/9.4001**, and some traffic to **self**.



NOTE: You can dynamically associate OIF maps with IGMP interfaces using dynamic profiles.

- **passive**—Defines either IGMP or MLD to use passive mode.

The OIF map interface should not typically pass IGMP or MLD control traffic and should be configured as passive. However, the OIF map implementation does support running IGMP or MLD on an interface (control and data) in addition to mapping data streams to the same interface. In this case, you should configure IGMP or MLD normally (that is, not in passive mode) on the mapped interface. In this example, the OIF map interfaces (**ge-2/3/9.4000** and **ge-2/3/9.4001**) are configured as MLD passive.

By default, specifying the **passive** statement means that no general queries, group-specific queries, or group-source-specific queries are sent over the interface and that all received control traffic is ignored by the interface. However, you can selectively activate up to two out of the three available options for the **passive** statement while keeping the other functions passive (inactive).

These options include the following:

- **send-general-query**—When specified, the interface sends general queries.
- **send-group-query**—When specified, the interface sends group-specific and group-source-specific queries.
- **allow-receive**—When specified, the interface receives control traffic.

Topology

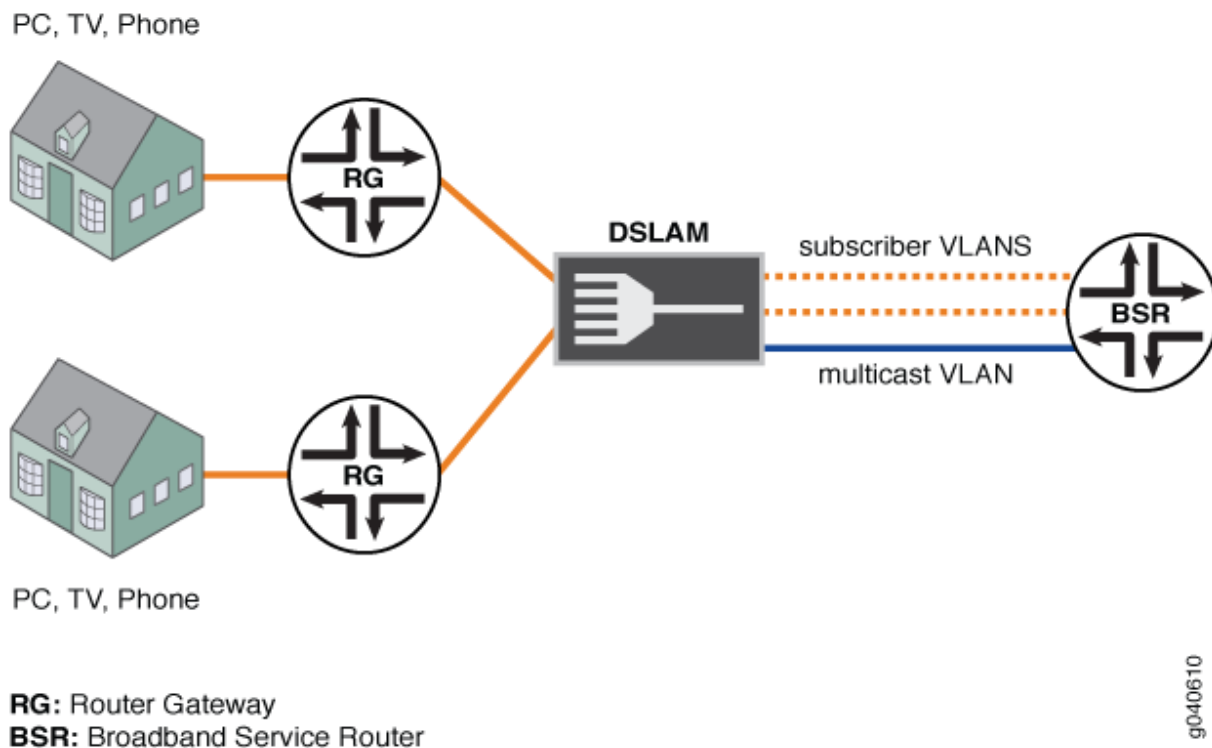
Figure 147 on page 1295 shows the scenario.

In both approaches, if multiple customers request the same multicast stream, the BSR sends one copy of the stream on the shared multicast VLAN interface. The DSLAM receives the multicast stream from the BSR on the shared interface and performs the necessary replication to the customers.

In the first approach (reverse OIF mapping), the DSLAM uses the per-customer subscriber VLANs for unicast data only. IGMP and MLD join and leave requests are sent on the multicast VLAN.

In the second approach (OIF mapping), the DSLAM uses the per-customer subscriber VLANs for unicast data and for IGMP and MLD join and leave requests. The multicast VLAN is used only for multicast streams, not for join and leave requests.

Figure 147: Multicast with Subscriber VLANs



Configuration

IN THIS SECTION

- [Configuring a Reverse OIF Map | 1296](#)
- [Configuring an OIF Map | 1301](#)

Configuring a Reverse OIF Map

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. .

```

set class-of-service traffic-control-profiles tcp-ifl shaping-rate 20m
set class-of-service interfaces ge-2/2/0 shaping-rate 240m
set class-of-service interfaces ge-2/2/0 unit 50 output-traffic-control-profile tcp-ifl
set class-of-service interfaces ge-2/2/0 unit 51 output-traffic-control-profile tcp-ifl
set interfaces ge-2/0/0 unit 0 family inet address 30.0.0.2/24
set interfaces ge-2/2/0 hierarchical-scheduler
set interfaces ge-2/2/0 vlan-tagging
set interfaces ge-2/2/0 unit 10 vlan-id 10
set interfaces ge-2/2/0 unit 10 family inet address 40.0.0.2/24
set interfaces ge-2/2/0 unit 50 vlan-id 50
set interfaces ge-2/2/0 unit 50 family inet address 50.0.0.2/24
set interfaces ge-2/2/0 unit 51 vlan-id 51
set interfaces ge-2/2/0 unit 51 family inet address 50.0.1.2/24
set policy-options policy-statement all-mcast-groups from source-address-filter 30.0.0.0/8
orlonger
set policy-options policy-statement all-mcast-groups then accept
set protocols igmp interface all
set protocols igmp interface fxp0.0 disable
set protocols pim rp local address 20.0.0.2
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface ge-2/2/0.10 disable
set routing-options multicast flow-map map1 policy all-mcast-groups

```

```

set routing-options multicast flow-map map1 bandwidth 10m
set routing-options multicast flow-map map1 bandwidth adaptive
set routing-options multicast interface ge-2/2/0.10 maximum-bandwidth 500m
set routing-options multicast interface ge-2/2/0.10 reverse-oif-mapping
set routing-options multicast interface ge-2/2/0.10 subscriber-leave-timer 20

```

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 [Junos OS CLI User Guide](#).

To configure reverse OIF mapping:

1. Configure a logical interface for unicast data traffic.

```

[edit interfaces ge-2/0/0]
user@host# set unit 0 family inet address 30.0.0.2/24

```

2. Configure a logical interface for subscriber control traffic.

```

[edit interfaces ge-2/2/0]
user@host# set hierarchical-scheduler
user@host# set vlan-tagging
user@host# set unit 10 vlan-id 10
user@host# set unit 10 family inet address 40.0.0.2/24

```

3. Configure two logical interfaces on which QoS adjustments are made.

```

[edit interfaces ge-2/2/0]
user@host# set unit 50 vlan-id 50
user@host# set unit 50 family inet address 50.0.0.2/24
user@host# set unit 51 vlan-id 51
user@host# set unit 51 family inet address 50.0.1.2/24

```

4. Configure a policy.

```
[edit policy-options policy-statement all-mcast-groups]
user@host# set from source-address-filter 30.0.0.0/8 orlonger
user@host# set then accept
```

5. Enable a flow map that references the policy.

```
[edit routing-options multicast]
user@host# set flow-map map1 policy all-mcast-groups
user@host# set flow-map map1 bandwidth 10m adaptive
```

6. Enable OIF mapping on the logical interface that receives subscriber control traffic.

```
[edit routing-options multicast]
user@host# set interface ge-2/2/0.10 maximum-bandwidth 500m
user@host# set interface ge-2/2/0.10 reverse-oif-mapping
user@host# set interface ge-2/2/0.10 subscriber-leave-timer 20
```

7. Configure PIM and IGMP.

```
[edit protocols]
user@host# set igmp interface all
user@host# set igmp interface fxp0.0 disable
user@host# set pim rp local address 20.0.0.2
user@host# set pim interface all
user@host# set pim interface fxp0.0 disable
user@host# set pim interface ge-2/2/0.10 disable
```

8. Configure the hierarchical scheduler by configuring a shaping rate for the physical interface and a slower shaping rate for the logical interfaces on which QoS adjustments are made.

```
[edit class-of-service interfaces ge-2/2/0]
user@host# set shaping-rate 240m
user@host# set unit 50 output-traffic-control-profile tcp-ifl
user@host# set unit 51 output-traffic-control-profile tcp-ifl
```

```
[edit class-of-service traffic-control-profiles tcp-30m-no-smap]
user@host# set shaping-rate 20m
```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service**, **show interfaces**, **show policy-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@host# show class-of-service
traffic-control-profiles {
  tcp-ifl {
    shaping-rate 20m;
  }
}
interfaces {
  ge-2/2/0 {
    shaping-rate 240m;
    unit 50 {
      output-traffic-control-profile tcp-ifl;
    }
    unit 51 {
      output-traffic-control-profile tcp-ifl;
    }
  }
}
```

```
user@host# show interfaces
ge-2/0/0 {
  unit 0 {
    family inet {
      address 30.0.0.2/24;
    }
  }
}
ge-2/2/0 {
  hierarchical-scheduler;
  vlan-tagging;
  unit 10 {
```

```
    vlan-id 10;
    family inet {
        address 40.0.0.2/24;
    }
}
unit 50 {
    vlan-id 50;
    family inet {
        address 50.0.0.2/24;
    }
}
unit 51 {
    vlan-id 51;
    family inet {
        address 50.0.1.2/24;
    }
}
}
```

```
user@host# show policy-options
policy-statement all-mcast-groups {
    from {
        source-address-filter 30.0.0.0/8 orlonger;
    }
    then accept;
}
```

```
user@host# show protocols
igmp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
pim {
    rp {
        local {
            address 20.0.0.2;
        }
    }
}
```

```

interface all;
interface fxp0.0 {
    disable;
}
interface ge-2/2/0.10 {
    disable;
}
}

```

```

user@host# show routing-options
multicast {
    flow-map map1 {
        policy all-mcast-groups;
        bandwidth 10m adaptive;
    }
    interface ge-2/2/0.10 {
        maximum-bandwidth 500m;
        reverse-oif-mapping;
        subscriber-leave-timer 20;
    }
}

```

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

Configuring an OIF Map

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.

```

set interfaces ge-2/3/8 unit 0 family inet6 address C300:0101::/24
set interfaces ge-2/3/9 vlan-tagging
set interfaces ge-2/3/9 unit 1 vlan-id 1
set interfaces ge-2/3/9 unit 1 family inet6 address C400:0101::/24
set interfaces ge-2/3/9 unit 2 vlan-id 2
set interfaces ge-2/3/9 unit 2 family inet6 address C400:0201::/24
set interfaces ge-2/3/9 unit 4000 vlan-id 4000
set interfaces ge-2/3/9 unit 4000 family inet6 address C40F:A001::/24

```

```

set interfaces ge-2/3/9 unit 4001 vlan-id 4001
set interfaces ge-2/3/9 unit 4001 family inet6 address C40F:A101::/24
set policy-options policy-statement g539-v6 term g539-4000 from route-filter FF05:0101:0000::/39
orlonger
set policy-options policy-statement g539-v6 term g539-4000 then map-to-interface ge-2/3/9.4000
set policy-options policy-statement g539-v6 term g539-4000 then accept
set policy-options policy-statement g539-v6 term g539-4001 from route-filter FF05:0101:0200::/39
orlonger
set policy-options policy-statement g539-v6 term g539-4001 then map-to-interface ge-2/3/9.4001
set policy-options policy-statement g539-v6 term g539-4001 then accept
set policy-options policy-statement g539-v6 term self from route-filter FF05:0101:0700::/40
orlonger
set policy-options policy-statement g539-v6 term self then map-to-interface self
set policy-options policy-statement g539-v6 term self then accept
set policy-options policy-statement g539-v6-all term g539 from route-filter 0::/0 orlonger
set policy-options policy-statement g539-v6-all term g539 then map-to-interface ge-2/3/9.4000
set policy-options policy-statement g539-v6-all term g539 then accept
set protocols mld interface fxp0.0 disable
set protocols mld interface ge-2/3/9.4000 passive
set protocols mld interface ge-2/3/9.4001 passive
set protocols mld interface ge-2/3/9.1 version 1
set protocols mld interface ge-2/3/9.1 oif-map g539-v6
set protocols mld interface ge-2/3/9.2 version 2
set protocols mld interface ge-2/3/9.2 oif-map g539-v6
set protocols pim rp local address 20.0.0.4
set protocols pim rp local family inet6 address C000::1
set protocols pim interface ge-2/3/8.0 mode sparse
set protocols pim interface ge-2/3/8.0 version 2
set routing-options multicast interface ge-2/3/9.1 no-qos-adjust
set routing-options multicast interface ge-2/3/9.2 no-qos-adjust

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the [Junos OS CLI User Guide](#).

To configure reverse OIF mapping:

1. Configure a logical interface for unicast data traffic.

```
[edit interfaces ge-2/3/8 ]
user@host# set unit 0 family inet6 address C300:0101::/24
```

2. Configure logical interfaces for subscriber VLANs.

```
[edit interfaces ge-2/3/9]
user@host# set vlan-tagging
user@host# set unit 1 vlan-id 1
user@host# set unit 1 family inet6 address C400:0101::/24
user@host# set unit 2 vlan-id 2
user@host# set unit 2 family inet6 address C400:0201::/24 lo0 unit 0 family inet6 address
C000::1/128
user@host# set unit 2 family inet6 address C400:0201::/24
```

3. Configure two map-to logical interfaces.

```
[edit interfaces ge-2/2/0]
user@host# set unit 4000 vlan-id 4000
user@host# set unit 4000 family inet6 address C40F:A001::/24
user@host# set unit 4001 vlan-id 4001
user@host# set unit 4001 family inet6 address C40F:A101::/24
```

4. Configure the OIF map.

```
[edit policy-options policy-statement g539-v6]
user@host# set term g539-4000 from route-filter FF05:0101:0000::/39 orlonger
user@host# set then map-to-interface ge-2/3/9.4000
user@host# set then accept
user@host# set term g539-4001 from route-filter FF05:0101:0200::/39 orlonger
user@host# set then map-to-interface ge-2/3/9.4001
user@host# set then accept
user@host# set term self from route-filter FF05:0101:0700::/40 orlonger
user@host# set then map-to-interface self
user@host# set then accept
[edit policy-options policy-statement g539-v6-all]
user@host# set term g539 from route-filter 0::/0 orlonger
```

```

user@host# set then map-to-interface ge-2/3/9.4000
user@host# set then accept

```

5. Disable QoS adjustment on the subscriber VLANs.

```

[edit routing-options multicast]
user@host# set interface ge-2/3/9.1 no-qos-adjust
user@host# set interface ge-2/3/9.2 no-qos-adjust

```

6. Configure PIM and MLD. Point the MLD subscriber VLANs to the OIF map.

```

[edit protocols]
user@host# set pim rp local address 20.0.0.4
user@host# set pim rp local family inet6 address C000::1 #C000::1 is the address of lo0
user@host# set pim interface ge-2/3/8.0 mode sparse
user@host# set pim interface ge-2/3/8.0 version 2
user@host# set mld interface fxp0.0 disable
user@host# set interface ge-2/3/9.4000 passive
user@host# set interface ge-2/3/9.4001 passive
user@host# set interface ge-2/3/9.1 version 1
user@host# set interface ge-2/3/9.1 oif-map g539-v6
user@host# set interface ge-2/3/9.2 version 2
user@host# set interface ge-2/3/9.2 oif-map g539-v6

```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-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@host# show interfaces
ge-2/3/8 {
  unit 0 {
    family inet6 {
      address C300:0101::/24;
    }
  }
}
ge-2/3/9 {

```

```
vlan-tagging;
unit 1 {
    vlan-id 1;
    family inet6 {
        address C400:0101::/24;
    }
}
unit 2 {
    vlan-id 2;
    family inet6 {
        address C400:0201::/24;
    }
}
unit 4000 {
    vlan-id 4000;
    family inet6 {
        address C40F:A001::/24;
    }
}
unit 4001 {
    vlan-id 4001;
    family inet6 {
        address C40F:A101::/24;
    }
}
}
```

```
user@host# show policy-options
policy-statement g539-v6 {
    term g539-4000 {
        from {
            route-filter FF05:0101:0000::/39 orlonger;
        }
        then {
            map-to-interface ge-2/3/9.4000;
            accept;
        }
    }
    term g539-4001 {
        from {
            route-filter FF05:0101:0200::/39 orlonger;
```

```
    }
    then {
        map-to-interface ge-2/3/9.4001;
        accept;
    }
}
term self {
    from {
        route-filter FF05:0101:0700::/40 orlonger;
    }
    then {
        map-to-interface self;
        accept;
    }
}
}
policy-statement g539-v6-all {
    term g539 {
        from {
            route-filter 0::/0 orlonger;
        }
        then {
            map-to-interface ge-2/3/9.4000;
            accept;
        }
    }
}
```

```
user@host# show protocols
mld {
    interface fxp0.0 {
        disable;
    }
    interface ge-2/3/9.4000 {
        passive;
    }
    interface ge-2/3/9.4001 {
        passive;
    }
    interface ge-2/3/9.1 {
        version 1;
    }
}
```

```
        oif-map g539-v6;
    }
    interface ge-2/3/9.2 {
        version 2;
        oif-map g539-v6;
    }
}
pim {
    rp {
        local {
            address 20.0.0.4;
            family inet6 {
                address C000::1;
            }
        }
    }
}
interface ge-2/3/8.0 {
    mode sparse;
    version 2;
}
}
```

```
user@host# show routing-options
multicast {
    interface ge-2/3/9.1 no-qos-adjust;
    interface ge-2/3/9.2 no-qos-adjust;
}
```

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

Verification

To verify the configuration, run the following commands:

- **show igmp statistics**
- **show class-of-service interface**
- **show interfaces statistics**
- **show mld statistics**
- **show multicast interface**

- `show policy`

SEE ALSO

[Example: Configuring a Multicast Flow Map | 1321](#)

Configuring Multicast Routing over IP Demux Interfaces

In a subscriber management network, fields in packets sent from IP demux interfaces are intended to correspond to a specific client that resides on the other side of an aggregation device (for example, a Multiservice Access Node [MSAN]). However, packets sent from a Broadband Services Router (BSR) to an MSAN do not identify the demux interface. Once it obtains a packet, it is up to the MSAN device to determine which client receives the packet.

Depending on the intelligence of the MSAN device, determining which client receives the packet can occur in an inefficient manner. For example, when it receives IGMP control traffic, an MSAN might forward the control traffic to all clients instead of the one intended client. In addition, once a data stream destination is established, though an MSAN can use IGMP snooping to determine which hosts reside in a particular group and limit data streams to only that group, the MSAN still must send multiple copies of the data stream to each group member, even if that data stream is intended for only one client in the group.

Various multicast features, when combined, enable you to avoid the inefficiencies mentioned above. These features include the following:

- The ability to configure the IP demux interface **family** statement to use **inet** for either the numbered or unnumbered primary interface.
- The ability to configure IGMP on the primary interface to send general queries for all clients. The demux configuration prevents the primary IGMP interface from receiving any client IGMP control packets. Instead, all IGMP control packets go to the demux interfaces. However, to guarantee that no joins occur on the primary interface:
 - For static IGMP interfaces—Include the **passive send-general-query** statement in the IGMP configuration at the `[edit protocols igmp interface interface-name]` hierarchy level.
 - For dynamic IGMP demux interfaces—Include the **passive send-general-query** statement at the `[edit dynamic-profiles profile-name protocols igmp interface interface-name]` hierarchy level.
- The ability to map all multicast groups to the primary interface as follows:
 - For static IGMP interfaces—Include the **oif-map** statement at the `[edit protocols igmp interface interface-name]` hierarchy level.
 - For dynamic IGMP demux interfaces—Include the **oif-map** statement at the `[edit dynamic-profiles profile-name protocols igmp interface interface-name]` hierarchy level.

Using the **oif-map** statement, you can map the same IGMP group to the same output interface and send only one copy of the multicast stream from the interface.

- The ability to configure IGMP on each demux interface. To prevent duplicate general queries:
 - For static IGMP interfaces—Include the **passive allow-receive send-group-query** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.
 - For dynamic demux interfaces—Include the **passive allow-receive send-group-query** statement at the **[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]** hierarchy level.



NOTE: To send only one copy of each group, regardless of how many customers join, use the **oif-map** statement as previously mentioned.

SEE ALSO

Example: Configuring Multicast with Subscriber VLANs

[Junos OS Subscriber Management and Services Library](#)

Classifying Packets by Egress Interface

For Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers with the Intelligent Queuing (IQ), IQ2, Enhanced IQ (IQE), Multiservices link services intelligent queuing (LSQ) interfaces, or ATM2 PICs, you can classify unicast and multicast packets based on the egress interface. For unicast traffic, you can also use a multifield filter, but only egress interface classification applies to multicast traffic as well as unicast traffic. If you configure egress classification of an interface, you cannot perform Differentiated Services code point (DSCP) rewrites on the interface. By default, the system does not perform any classification based on the egress interface.

On an MX Series router that contains MPCs and MS-DPCs, multicast packets are dropped on the router and not processed properly if the router contains MLPPP LSQ logical interfaces that function as multicast receivers and if the network services mode is configured as enhanced IP mode on the router. This behavior is expected with LSQ interfaces in conjunction with enhanced IP mode. In such a scenario, if enhanced IP mode is not configured, multicasting works correctly. However, if the router contains redundant LSQ interfaces and enhanced IP network services mode configured with FIB localization, multicast works properly.

To enable packet classification by the egress interface, you first configure a forwarding class map and one or more queue numbers for the egress interface at the [edit class-of-service forwarding-class-map *forwarding-class-map-name*] hierarchy level:

```
[edit class-of-service]
forwarding-classes-interface-specific forwarding-class-map-name {
  class class-name queue-num queue-number [ restricted-queue queue-number ];
}
```

For T Series routers that are restricted to only four queues, you can control the queue assignment with the `restricted-queue` option, or you can allow the system to automatically determine the queue in a modular fashion. For example, a map assigning packets to queue 6 would map to queue 2 on a four-queue system.



NOTE: If you configure an output forwarding class map associating a forwarding class with a queue number, this map is not supported on multiservices link services intelligent queuing (lsq-) interfaces.

Once the forwarding class map has been configured, you apply the map to the logical interface by using the `output-forwarding-class-map` statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
output-forwarding-class-map forwarding-class-map-name;
```

All parameters relating to the queues and forwarding class must be configured as well. For more information about configuring forwarding classes and queues, see [Configuring a Custom Forwarding Class for Each Queue](#).

This example shows how to configure an interface-specific forwarding-class map named FCMAP1 that restricts queues 5 and 6 to different queues on four-queue systems and then applies FCMAP1 to unit 0 of interface `ge-6/0/0`:

```
[edit class-of-service]
forwarding-class-map FCMAP1 {
  class FC1 queue-num 6 restricted-queue 3;
  class FC2 queue-num 5 restricted-queue 2;
  class FC3 queue-num 3;
  class FC4 queue-num 0;
  class FC3 queue-num 0;
```



```

    class FC4 queue-num 1;
}

[edit class-of-service]
interfaces {
    ge-6/0/0 unit 0 {
        output-forwarding-class-map FCMAP1;
    }
}

```

Note that without the `restricted-queue` option in FCMAP1, the example would assign FC1 and FC2 to queues 2 and 1, respectively, on a system restricted to four queues.

Use the `show class-of-service forwarding-class forwarding-class-map-name` command to display the forwarding-class map queue configuration:

```
user@host> show class-of-service forwarding-class FCMAP2
```

Forwarding class	ID	Queue	Restricted queue
FC1	0	6	3
FC2	1	5	2
FC3	2	3	3
FC4	3	0	0
FC5	4	0	0
FC6	5	1	1
FC7	6	6	2
FC8	7	7	3

Use the `show class-of-service interface interface-name` command to display the forwarding-class maps (and other information) assigned to a logical interface:

```
user@host> show class-of-service interface ge-6/0/0
```

```

Physical interface: ge-6/0/0, Index: 128
Queues supported: 8, Queues in use: 8
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4

```

```
Logical interface: ge-6/0/0.0, Index: 67
```

Object	Name	Type	Index
Scheduler-map	sch-map1	Output	6998

Scheduler-map	sch-map1	Input	6998
Classifier	dot1p	ieee8021p	4906
forwarding-class-map	FCMAP1	Output	1221
Logical interface: ge-6/0/0.1, Index 68			
Object	Name	Type	Index
Scheduler-map	<default>	Output	2
Scheduler-map	<default>	Input	3
Logical interface: ge-6/0/0.32767, Index 69			
Object	Name	Type	Index
Scheduler-map	<default>	Output	2
Scheduler-map	<default>	Input	3

Recycle Bandwidth Management

IN THIS SECTION

- [Default Recycle Bandwidth Mode | 1313](#)
- [Configurable Recycle Bandwidth Mode | 1314](#)

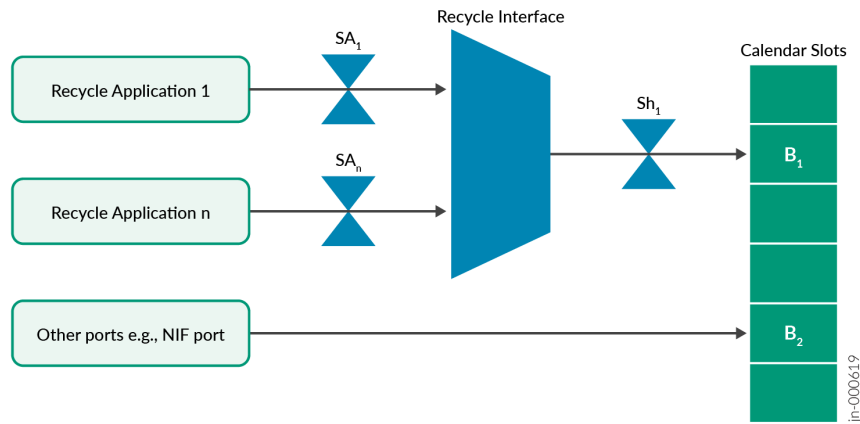
ACX routers use the recycle interface to loopback or recirculate traffic from the egress interface to the ingress for additional processing. This is needed for applications that require additional processing.

The recycle interface is an internal channelized interface that has an egress shaper that limits the flow rate of outbound traffic. By default, the recycle interface bandwidth is based on chassis oversubscription and platform specific FPC or interface configuration.

Applications that use the recycle mechanism, configure channels or virtual ports based on need, which in turn goes through an application shaper similar to the egress shaper on the recycle interface. The recycle interface can support a maximum of 256 channels or virtual logical ports. All the recycle channels operate at the same priority. The recycled traffic is then forwarded using the calendar slot based on the weight assigned to it. Other types of interfaces like the NIF interface forward traffic by occupying other slots on the calendar proportional to its bandwidth.

Figure 1 illustrates the the recycle mechanism and its components.

Figure 148: The recycle mechanism



- Sh_1 – Recycle interface shaper
- SA_1 – Application #1 shaper over recycle channel 1
- SA_n – Application #n shaper over recycle channel n (max 256)
- B_1 – Calendar slot weight for traffic from the recycle interface. This reflects the calendar bandwidth value allocated to the recycle interface.
- B_2 – Calendar slot weight for traffic from interfaces other than the recycle interface. For example the NIF interface.

The recycle mechanism has two modes of operation:

- Default Recycle Bandwidth Mode
- Configurable Recycle Bandwidth Mode

Default Recycle Bandwidth Mode

Like the name suggests, the default recycle bandwidth mode is enabled by default and user configuration is not necessary. The recycle interface is allocated a set portion of the total calendar bandwidth. This recycle bandwidth is guaranteed and is therefore stable in times of traffic congestion. The recycle applications share this bandwidth on a best effort basis, which means there is no guaranteed bandwidth per application.

Benefits of default recycle bandwidth include:

- Efficient utilization of recycle bandwidth. In case no recycle applications are running, the unused bandwidth can be shared among running applications from other interfaces.

- Efficient utilization of chip bandwidth. In case the other interfaces other than the recycle interface are not carrying traffic, the unused bandwidth can be utilized by the recycle applications.

Configurable Recycle Bandwidth Mode

By configuring the recycle interface based on profiles, application bandwidth becomes manageable. You can guarantee allocation of recycle bandwidth for an application by defining it in a profile.

Benefits of configurable recycle bandwidth include:

- Guaranteed bandwidth allocation for a defined recycle application.
- Flexibility to change bandwidth allocation, which helps you to prioritize application bandwidth.

Example: Configuring Recycle Bandwidth

SUMMARY

This example shows how to manage recycle bandwidth per application.

IN THIS SECTION

- [Overview | 1314](#)
- [Configure Recycle Interface Bandwidth | 1315](#)
- [Verification | 1315](#)

Overview

Applications 1, 2, 3, and 4 are recycle applications that use the recycle interface. Table 1 shows our bandwidth requirements. In this case, we want guaranteed 10% allocation of the recycle interface bandwidth for application 1 and 20% for application 2. Applications 3 and 4 are not priority and no guarantee is necessary.

Table 37: Bandwidth allocation per recycle application

Recycle Application	Required recycle bandwidth value in percentage
Application 1	10%
Application 2	20%
Application 3	undefined

Table 37: Bandwidth allocation per recycle application (*Continued*)

Recycle Application	Required recycle bandwidth value in percentage
Application 4	undefined

Configure Recycle Interface Bandwidth

1. set system packet-forwarding-options recycle-bandwidth profile *profile1*
2. set system packet-forwarding-options recycle-bandwidth-profile *profile1 application1 10 application2 20*

Verification

IN THIS SECTION

- [Verifying recycle bandwidth per application | 1315](#)

Command	Verification Task
show system packet-forwarding-options recycle-bandwidth-profile	From operational mode, run the show system packet-forwarding-options recycle-bandwidth-profile command.

Verifying recycle bandwidth per application

Purpose

Verifying recycle bandwidth allocation per application.

Action

```
user@router> show system packet-forwarding-options recycle-bandwidth-profile
```

Recycle Interface details:

```

Total Bandwidth :    300 Gbps/Core
PFE : 0
Application-Name   | Core   | Bandwidth(Kbps) | Port   | VoQ   | Profile
-----
application1       | 0      | 300000000       | 234    | 1336  | profile1
application2       | 0      | 600000000       | 242    | 1400  | profile1
application3       | 0      | 1050000000      | 243    | 1408  | profile1
application4       | 0      | 1050000000      | 245    | 1424  | profile1
PFE : 1
Application-Name   | Core   | Bandwidth(Kbps) | Port   | VoQ   | Profile
-----
application1       | 0      | 300000000       | 234    | 2880  | profile1
application2       | 0      | 600000000       | 242    | 2944  | profile1
application3       | 0      | 1050000000      | 243    | 2952  | profile1
application4       | 0      | 1050000000      | 245    | 2968  | profile1

```

Meaning

The output shows the recycle bandwidth allocation per application as configured in the section earlier.

Table 38: Recycle Bandwidth Allotment

Application	Percentage configured	Result
application1	10	300 Gbps x 10% = 30 Gbps
application2	20	300 Gbps x 20% = 60 Gbps
application3	not configured	Unutilized bandwidth/number of unconfigured recycle applications. (300 x 70%) / 2 = 105 Gbps
application4	not configured	(300 x 70%) / 2 = 105 Gbps

SEE ALSO

| *recycle-bandwidth-profiles*

RELATED DOCUMENTATION

[Examples: Configuring Administrative Scoping | 1273](#)

[Examples: Configuring the Multicast Forwarding Cache | 1317](#)

Examples: Configuring the Multicast Forwarding Cache

IN THIS SECTION

- [Understanding the Multicast Forwarding Cache | 1317](#)
- [Example: Configuring the Multicast Forwarding Cache | 1317](#)
- [Example: Configuring a Multicast Flow Map | 1321](#)

Understanding the Multicast Forwarding Cache

IP multicast protocols can create numerous entries in the multicast forwarding cache. If the forwarding cache fills up with entries that prevent the addition of higher-priority entries, applications and protocols might not function properly. You can manage the multicast forwarding cache properties by limiting the size of the cache and by controlling the length of time that entries remain in the cache. By managing timeout values, you can give preference to more important forwarding cache entries while removing other less important entries.

Example: Configuring the Multicast Forwarding Cache

IN THIS SECTION

- [Requirements | 1318](#)
- [Overview | 1318](#)
- [Configuration | 1319](#)
- [Verification | 1321](#)

When a routing device receives multicast traffic, it places the (S,G) route information in the multicast forwarding cache, **inet.1**. This example shows how to configure multicast forwarding cache limits to prevent the cache from filling up with entries.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure a multicast protocol. This feature works with the following multicast protocols:
 - DVMRP
 - PIM-DM
 - PIM-SM
 - PIM-SSM

Overview

IN THIS SECTION

- [Topology | 1319](#)

This example includes the following statements:

- **forwarding-cache**—Specifies how forwarding entries are aged out and how the number of entries is controlled.
- **timeout**—Specifies an idle period after which entries are aged out and removed from **inet.1**. You can specify a timeout in the range from 1 through 720 minutes.
- **threshold**—Enables you to specify threshold values on the forwarding cache to suppress (suspend) entries from being added when the cache entries reach a certain maximum and begin adding entries to the cache when the number falls to another threshold value. By default, no threshold values are enabled on the routing device.

The suppress threshold suspends the addition of new multicast forwarding cache entries. If you do not specify a suppress value, multicast forwarding cache entries are created as necessary. If you specify a suppress threshold, you can optionally specify a reuse threshold, which sets the point at which the device resumes adding new multicast forwarding cache entries. During suspension, forwarding cache entries time out. After a certain number of entries time out, the reuse threshold is reached, and new entries are added. The range for both thresholds is from 1 through 200,000. If configured, the reuse value must be less than the suppression value. If you do not specify a reuse value, the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.

Topology

Configuration

IN THIS SECTION

- [Procedure | 1319](#)
- [Results | 1320](#)

Procedure

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.

```
set routing-options multicast forwarding-cache threshold suppress 150000
set routing-options multicast forwarding-cache threshold reuse 34
set routing-options multicast forwarding-cache timeout 60
```

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 [Junos OS CLI User Guide](#).

To configure the multicast forwarding cache:

1. Configure the maximum size of the forwarding cache.

```
[edit routing-options multicast forwarding-cache]
user@host# set threshold suppress 150000
```

2. Configure the amount of time (in minutes) entries can remain idle before being removed.

```
[edit routing-options multicast forwarding-cache]
user@host# set timeout 60
```

3. Configure the size of the forwarding cache when suppression stops and new entries can be added.

```
[edit routing-options multicast forwarding-cache]
user@host# set threshold reuse 70000
```

Results

Confirm your configuration by entering the **show routing-options** command.

```
user@host# show routing-options
multicast {
  forwarding-cache {
    threshold {
      suppress 150000;
      reuse 70000;
    }
    timeout 60;
  }
}
```

Verification

To verify the configuration, run the **show multicast route extensive** command.

```
user@host> show multicast route extensive
Family: INET
Group: 232.0.0.1
  Source: 11.11.11.11/32
  Upstream interface: fe-0/2/0.200
  Downstream interface list:
    fe-0/2/1.210
  Downstream interface list rejected by CAC:
    fe-0/2/1.220
  Session description: Source specific multicast
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 337
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 60 minutes
  Wrong incoming interface notifications: 0
```

SEE ALSO

[Bandwidth Management and Source Redundancy | 1285](#)

[Understanding Bandwidth Management for Multicast | 1284](#)

Example: Configuring a Multicast Flow Map

IN THIS SECTION

- [Requirements | 1322](#)
- [Overview | 1322](#)
- [Configuration | 1323](#)
- [Verification | 1326](#)

This example shows how to configure a flow map to prevent certain forwarding cache entries from aging out, thus allowing for faster failover from one source to another. Flow maps enable you to configure bandwidth variables and multicast forwarding cache timeout values for entries defined by the flow map policy.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Library for Routing Devices](#).
- Configure a multicast protocol. This feature works with the following multicast protocols:
 - DVMRP
 - PIM-DM
 - PIM-SM
 - PIM-SSM

Overview

Flow maps are typically used for fast multicast source failover when there are multiple sources for the same group. For example, when one video source is actively sending the traffic, the forwarding states for other video sources are timed out after a few minutes. Later, when a new source starts sending the traffic again, it takes time to install a new forwarding state for the new source if the forwarding state is not already there. This switchover delay is worsened when there are many video streams. Using flow maps with longer timeout values or permanent cache entries helps reduce this switchover delay.



NOTE: The permanent forwarding state must exist on all routing devices in the path for fast source switchover to function properly.

This example includes the following statements:

- **bandwidth**—Specifies the bandwidth for each flow that is defined by a flow map to ensure that an interface is not oversubscribed for multicast traffic. If adding one more flow would cause overall bandwidth to exceed the allowed bandwidth for the interface, the request is rejected. A rejected request means that traffic might not be delivered out of some or all of the expected outgoing interfaces. You can define the bandwidth associated with multicast flows that match a flow map by specifying a bandwidth in bits per second or by specifying that the bandwidth is measured and adaptively modified.

When you use the **adaptive** option, the bandwidth adjusts based on measurements made at 5-second intervals. The flow uses the maximum bandwidth value from the last 12 measured values (1 minute).

When you configure a bandwidth value with the **adaptive** option, the bandwidth value acts as the starting bandwidth for the flow. The bandwidth then changes based on subsequent measured bandwidth values. If you do not specify a bandwidth value with the **adaptive** option, the starting bandwidth defaults to 2 megabits per second (Mbps).

For example, the **bandwidth 2m adaptive** statement is equivalent to the **bandwidth adaptive** statement because they both use the same starting bandwidth (2 Mbps, the default). If the actual flow bandwidth is 4 Mbps, the measured flow bandwidth changes to 4 Mbps after reaching the first measuring point (5 seconds). However, if the actual flow bandwidth rate is 1 Mbps, the measured flow bandwidth remains at 2 Mbps for the first 12 measurement cycles (1 minute) and then changes to the measured 1 Mbps value.

- **flow-map**—Defines a flow map that controls the forwarding cache timeout of specified source and group addresses, controls the bandwidth for each flow, and specifies redundant sources. If a flow can match multiple flow maps, the first flow map applies.
- **forwarding-cache**—Enables you to configure the forwarding cache properties of entries defined by a flow map. You can specify a timeout of **never** to make the forwarding entries permanent, or you can specify a timeout in the range from 1 through 720 minutes. If you set the value to **never**, you can specify the **non-discard-entry-only** option to make an exception for entries that are in the pruned state. In other words, the **never non-discard-entry-only** statement allows entries in the pruned state to time out, while entries in the forwarding state never time out.
- **policy**—Specifies source and group addresses to which the flow map applies.
- **redundant-sources**—Specify redundant (backup) sources for flows identified by a flow map. Outbound interfaces that are admitted for one of the forwarding entries are automatically admitted for any other entries identified by the redundant source configuration. In the example that follows, the two forwarding entries, (10.11.11.11) and (10.11.11.12,) match the flow map defined for **flowMap1**. If an outbound interface is admitted for entry (10.11.11.11), it is also automatically admitted for entry (10.11.11.12) so one source or the other can send traffic at any time.

Configuration

IN THIS SECTION

- Procedure | [1324](#)
- Results | [1325](#)

Procedure

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.

```
set policy-options prefix-list permanentEntries1 232.1.1.0/24
set policy-options policy-statement policyForFlow1 from source-address-filter 11.11.11.11/32
exact
set policy-options policy-statement policyForFlow1 from prefix-list-filter permanentEntries1
orlonger
set policy-options policy-statement policyForFlow1 then accept
set routing-options multicast flow-map flowMap1 policy policyForFlow1
set routing-options multicast flow-map flowMap1 bandwidth 2m
set routing-options multicast flow-map flowMap1 bandwidth adaptive
set routing-options multicast flow-map flowMap1 redundant-sources 10.11.11.11
set routing-options multicast flow-map flowMap1 redundant-sources 10.11.11.12
set routing-options multicast flow-map flowMap1 forwarding-cache timeout never non-discard-entry-
only
```

Step-by-Step Procedure

Multicast flow maps enable you to manage a subset of multicast forwarding table entries. For example, you can specify that certain forwarding cache entries be permanent or have a different timeout value from other multicast flows that are not associated with the flow map policy.

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 [Junos OS CLI User Guide](#).

To configure a flow map:

1. Configure the flow map policy. This step creates a flow map policy called `policyForFlow1`. The policy statement matches the source address using the `source-address-filter` statement, and matches the group address using the `prefix-list-filter`. The addresses must match the configured policy for flow mapping to occur.

```
[edit policy-options]
user@host# set prefix-list permanentEntries1 232.1.1.0/24
user@host# set policy policyForFlow1 from source-address-filter 11.11.11.11/32 exact
```

```
user@host# set policy policyForFlow1 from prefix-list-filter permanentEntries1 orlonger
user@host# set policy policyForFlow1 then accept
```

2. Define a flow map, **flowMap1**, that references the flow map policy, **policyForFlow1**, we just created.

```
[edit routing-options]
user@host# set multicast flow-map flowMap1 policy policyForFlow1
```

3. Configure permanent forwarding entries (that is, entries that never time out), and enable entries in the pruned state to time out.

```
[edit routing-options]
user@host# set multicast flow-map flowMap1 forwarding-cache timeout never non-discard-entry-only
```

4. Configure the flow map bandwidth to be adaptive with a default starting bandwidth of 2 Mbps.

```
[edit routing-options]
user@host# set multicast flow-map flowMap1 bandwidth 2m adaptive
```

5. Specify backup sources.

```
[edit routing-options]
user@host# set multicast flow-map flowMap1 redundant-sources [ 10.11.11.11 10.11.11.12 ]
```

6. Commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** and **show routing-options** commands.

```
user@host# show policy-options
prefix-list permanentEntries1 {
  232.1.1.0/24;
```

```
}  
policy-statement policyForFlow1 {  
  from {  
    source-address-filter 11.11.11.11/32 exact;  
    prefix-list-filter permanentEntries1 orlonger;  
  }  
  then accept;  
}
```

```
user@host# show routing-options  
multicast {  
  flow-map flowMap1 {  
    policy policyForFlow1;  
    bandwidth 2m adaptive;  
    redundant-sources [ 10.11.11.11 10.11.11.12 ];  
    forwarding-cache {  
      timeout never non-discard-entry-only;  
    }  
  }  
}
```

Verification

To verify the configuration, run the following commands:

- **show multicast flow-map**
- **show multicast route extensive**

SEE ALSO

[Understanding Bandwidth Management for Multicast | 1284](#)

[Bandwidth Management and Source Redundancy | 1285](#)

RELATED DOCUMENTATION

[Examples: Configuring Administrative Scoping | 1273](#)

[Examples: Configuring Bandwidth Management | 1284](#)

Example: Configuring Ingress PE Redundancy

IN THIS SECTION

- [Understanding Ingress PE Redundancy | 1327](#)
- [Example: Configuring Ingress PE Redundancy | 1327](#)

Understanding Ingress PE Redundancy

In many network topologies, point-to-multipoint label-switched paths (LSPs) are used to distribute multicast traffic over a virtual private network (VPN). When traffic engineering is added to the provider edge (PE) routers, a popular deployment option has been to use traffic-engineered point-to-multipoint LSPs at the origin PE. In these network deployments, the PE is a single point of failure. Network operators have previously provided redundancy by broadcasting duplicate streams of multicast traffic from multiple PEs, a practice which at least doubles the bandwidth required for each stream.

Ingress PE redundancy eliminates the bandwidth duplication requirement by configuring one or more ingress PEs as a group. Within a group, one PE is designated as the primary PE and one or more others become backup PEs for the configured traffic stream. The solution depends on a full mesh of point-to-point (P2P) LSPs among the primary and backup PEs. Also, you must configure a full set of point-to-multipoint LSPs at the backup PEs, even though these point-to-multipoint LSPs at the backup PEs are not sending any traffic or using any bandwidth. The P2P LSPs are configured with bidirectional forwarding detection (BFD). When BFD detects a failure on the primary PE, a new designated forwarder is elected for the stream.

SEE ALSO

[MPLS Applications User Guide](#)

Example: Configuring Ingress PE Redundancy

IN THIS SECTION

- [Requirements | 1328](#)
- [Overview | 1328](#)
- [Configuration | 1330](#)

This example shows how to configure one PE as part of a backup PE group to enable ingress PE redundancy for multicast traffic streams.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure a full mesh of P2P LSPs between the PEs in the backup group.

Overview

Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution. When point-to-multipoint LSPs are used for multicast traffic, the PE device can become a single point of failure. One way to provide redundancy is by broadcasting duplicate streams from multiple PEs, thus doubling the bandwidth requirements for each stream. This feature implements redundancy between two or more PEs by designating a primary and one or more backup PEs for each configured stream. The solution depends on the configuration of a full mesh of P2P LSPs between the primary and backup PEs. These LSPs are configured with Bidirectional Forwarding Detection (BFD) running on top of them. BFD is used on the backup PEs to detect failure on the primary PE routing device and to elect a new designated forwarder for the stream.

A full mesh is required so that each member of the group can make an independent decision about the health of the other PEs and determine the designated forwarder for the group. The key concept in a backup PE group is that of a designated PE. A designated PE is a PE that forwards data on the static route. All other PEs in the backup PE group do not forward any data on the static route. This allows you to have one designated forwarder. If the designated forwarder fails, another PE takes over as the designated forwarder, thus allowing the traffic flow to continue uninterrupted.

Each PE in the backup PE group makes its own local decision regarding the designated forwarder. Thus, there is no inter-PE communication regarding designated forwarder. A PE computes the designated forwarder based on the IP address of all PEs and the connectivity status of other PEs. Connectivity status is determined based on the state of the BFD session on the P2P LSP to a PE.

A PE chosen is as the designated forwarder if it satisfies the following conditions:

- The PE is in the UP state. Either it is the local PE, or the BFD session on the P2P LSP to that PE is in the UP state.

- The PE has the lowest IP address among all PEs that are in the UP state.

Because all PEs have P2P LSPs to each other, each PE can determine the UP state of each other PE, and all PEs converge to the same designated forwarder.

If the designated forwarder PE fails, then all other PEs lose connectivity with the designated forwarder, and their BFD session ends. Consequently, other PEs then choose another designated forwarder. The new forwarder starts forwarding traffic. Thus, the traffic loss is limited to the failure detection time, which is the BFD session detection time.

When a PE that was the designated forwarder fails and then resumes operating, all other PEs recognize this fact, rerun the designated forwarder algorithm, and choose the PE as the designated forwarder. Consequently, the backup designated forwarder stops forwarding traffic. Thus, traffic switches back to the most eligible designated forwarder.

This example includes the following statements:

- **associate-backup-pe-groups**—Monitors the health of the routing device at the other end of the LSP. You can configure multiple backup PE groups that contain the same routing device's address. Failure of this LSP indicates to all of these groups that the destination PE routing device is down. So, the **associate-backup-pe-groups** statement is not tied to any specific group but applies to all groups that are monitoring the health of the LSP to the remote address.

If there are multiple LSPs with the **associate-backup-pe-groups** statement to the same destination PE, then the local routing device picks the first LSP to that PE for detection purposes.

We do not recommend configuring multiple LSPs to the same destination. If you do, make sure that the LSP parameters (for example, liveness detection) are similar to avoid false failure notification even when the remote PE is up.

- **backup-pe-group**—Configures ingress PE redundancy for multicast traffic streams.
- **bfd-liveness-detection**—Enables BFD for each LSP.
- **label-switched-path**—Configures an LSP. You must configure a full mesh of P2P LSPs between the primary and backup PEs.



NOTE: We recommend that you configure the P2P LSPs with fast reroute and node link protection so that link failures do not result in the LSP failure. For the purpose of PE redundancy, a failure in the P2P LSP is treated as a PE failure. Redundancy in the inter-PE path is also encouraged.

- **p2mp-lsp-next-hop**—Enables you to associate a backup PE group with a static route.

- **static**—Applies the backup group to a static route on the PE. This ensures that the static route is active (installed in the forwarding table) when the local PE is the designated forwarder for the configured backup PE group.

Configuration

IN THIS SECTION

● Procedure | [1330](#)

● Results | [1332](#)

Procedure

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.

```
set policy-options policy-statement no-rpf from route-filter 225.1.1.1/32 exact
set policy-options policy-statement no-rpf then reject
set protocols mpls label-switched-path backup_PE1 to 10.255.16.61
set protocols mpls label-switched-path backup_PE1 oam bfd-liveness-detection minimum-interval
500
set protocols mpls label-switched-path backup_PE1 oam bfd-liveness-detection multiplier 3
set protocols mpls label-switched-path backup_PE1 associate-backup-pe-groups
set protocols mpls label-switched-path dest1 to 10.255.16.57
set protocols mpls label-switched-path dest1 p2mp p2mp-lsp
set protocols mpls label-switched-path dest2 to 10.255.16.55
set protocols mpls label-switched-path dest2 p2mp p2mp-lsp
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set routing-options static route 10.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
set routing-options static route 10.1.1.1/32 backup-pe-group g1
set routing-options static route 225.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
set routing-options static route 225.1.1.1/32 backup-pe-group g1
set routing-options multicast rpf-check-policy no-rpf
set routing-options multicast interface fe-1/3/3.0 enable
```

```
set routing-options multicast backup-pe-group g1 backups 10.255.16.61
set routing-options multicast backup-pe-group g1 local-address 10.255.16.59
```

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 [Junos OS CLI User Guide](#).

To configure ingress PE redundancy:

1. Configure the multicast settings.

```
[edit routing-options multicast]
user@host# set rpf-check-policy no-rpf
user@host# set interface fe-1/3/3.0 enable
```

2. Configure the RPF policy.

```
[edit policy-options policy-statement no-rpf]
user@host# set from route-filter 225.1.1.1/32 exact
user@host# set then reject
```

3. Configure the backup PE group.

```
[edit routing-options multicast]
user@host# set backup-pe-group g1 backups 10.255.16.61
user@host# set backup-pe-group g1 local-address 10.255.16.59
```

4. Configure the static routes for the point-to-multipoint LSPs backup PE group.

```
[edit routing-options static]
user@host# set route 10.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
user@host# set route 10.1.1.1/32 backup-pe-group g1
user@host# set route 225.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
user@host# set route 225.1.1.1/32 backup-pe-group g1
```

5. Configure the MPLS interfaces.

```
[edit protocols mpls]
user@host# set interface all
user@host# set interface fxp0.0 disable
```

6. Configure the LSP to the redundant router.

```
[edit protocols mpls]
user@host# set label-switched-path backup_PE1 to 10.255.16.61
user@host# set label-switched-path backup_PE1 oam bfd-liveness-detection minimum-interval 500
user@host# set label-switched-path backup_PE1 oam bfd-liveness-detection multiplier 3
user@host# set label-switched-path backup_PE1 associate-backup-pe-groups
```

7. Configure LSPs to two traffic destinations.

```
[edit protocols mpls]
user@host# set label-switched-path dest1 to 10.255.16.57
user@host# set label-switched-path dest1 p2mp p2mp-lsp
user@host# set label-switched-path dest2 to 10.255.16.55
user@host# set label-switched-path dest2 p2mp p2mp-lsp
```

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

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy**, **show protocols**, and **show routing-options** commands.

```
user@host# show policy
policy-statement no-rpf {
  from {
    route-filter 225.1.1.1/32 exact;
  }
}
```

```
    then reject;
}
```

```
user@host# show protocols
mpls {
  label-switched-path backup_PE1 {
    to 10.255.16.61;
    oam {
      bfd-liveness-detection {
        minimum-interval 500;
        multiplier 3;
      }
    }
    associate-backup-pe-groups;
  }
  label-switched-path dest1 {
    to 10.255.16.57;
    p2mp p2mp-lsp;
  }
  label-switched-path dest2 {
    to 10.255.16.55;
    p2mp p2mp-lsp;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

```
user@host# show routing-options
static {
  route 10.1.1.1/32 {
    p2mp-lsp-next-hop p2mp-lsp;
    backup-pe-group g1;
  }
  route 225.1.1.1/32 {
    p2mp-lsp-next-hop p2mp-lsp;
    backup-pe-group g1;
  }
}
```

```
multicast {
  rpf-check-policy no-rpf;
  interface fe-1/3/3.0 enable;
  backup-pe-group g1 {
    backups 10.255.16.61;
    local-address 10.255.16.59;
  }
}
```

Verification

To verify the configuration, run the following commands:

- **show mpls lsp**
- **show multicast backup-pe-groups**
- **show multicast rpf**

SEE ALSO

Example: Configuring RPF Policies

RELATED DOCUMENTATION

[Examples: Configuring Administrative Scoping | 1273](#)

[Examples: Configuring Bandwidth Management | 1284](#)

[Examples: Configuring the Multicast Forwarding Cache | 1317](#)

Controller-Based BGP Multicast Signaling

IN THIS CHAPTER

- [Controller-Based BGP Multicast Signaling | 1335](#)

Controller-Based BGP Multicast Signaling

IN THIS SECTION

- [Benefits of Controller-Based BGP Multicast Signaling | 1335](#)
- [Transit Router Programming Based on PRPD | 1336](#)
- [Route Reflector Programming Based on PRPD and Traffic Transmission Using BGP | 1337](#)

Network controllers that are aware of the network topology and the events within the topology can calculate end-to-end explicit paths. This calculation results in optimum multicast trees between the source and the receivers. A network controller uses the programmable routing protocol daemon (PRPD) and BGP signaling to signal each router on the multicast tree to set up and program multicast forwarding states.

The controller uses PRPD APIs to program BGP multicast network layer reachability information (NLRI). All the routers are programmed either directly or through a route reflector. A route reflector also propagates the NLRI to all the routers on the trees. Based on the NLRI received, the routers set up appropriate forwarding states.

In the following sections of this topic we use *controller* to mean *network controller*.

Benefits of Controller-Based BGP Multicast Signaling

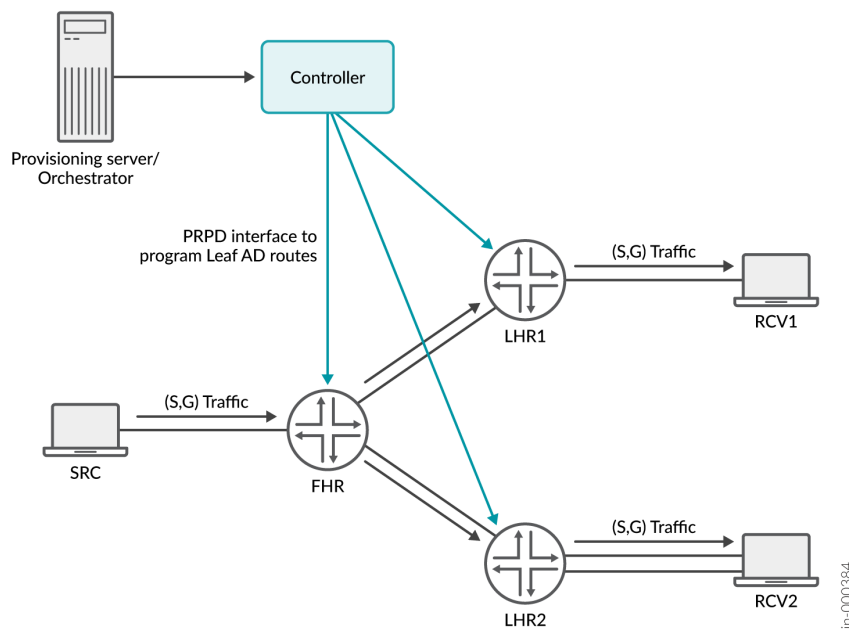
Traditional multicast deployments have inherent drawbacks. Controller-based BGP signaling eliminates drawbacks such as the following:

- In hop-by-hop multicast protocols such as Protocol Independent Multicast (PIM), among other limitations, the any-source multicast mode needs periodic refreshing and is complex to configure. PIM source-specific multicast and PIM-Port mode have similar drawbacks.
- Data centers that use BGP as their primary routing protocol avoid deploying multicast, because this deployment requires additional protocols and creates associated complexities.

Transit Router Programming Based on PRPD

With the help of a provisioning server or orchestrator, the controller is aware of the multicast flow information of the source, groups, and receivers. The controller calculates an optimum tree to forward the source traffic to all the interested receivers for a group.

Figure 149: Transit Router Programming Based on PRPD



After calculating the multicast tree, the controller uses the PRPD interface to program each transit router on the tree with the BGP multicast Leaf Auto Discovery (AD) NLRI. The Leaf AD route contains the source, group, and upstream information.

The Leaf AD NLRI also carries a tunnel encapsulation attribute (TEA) that has information about the downstream routers to which traffic needs to be replicated. TEA is a transitive BGP path attribute that

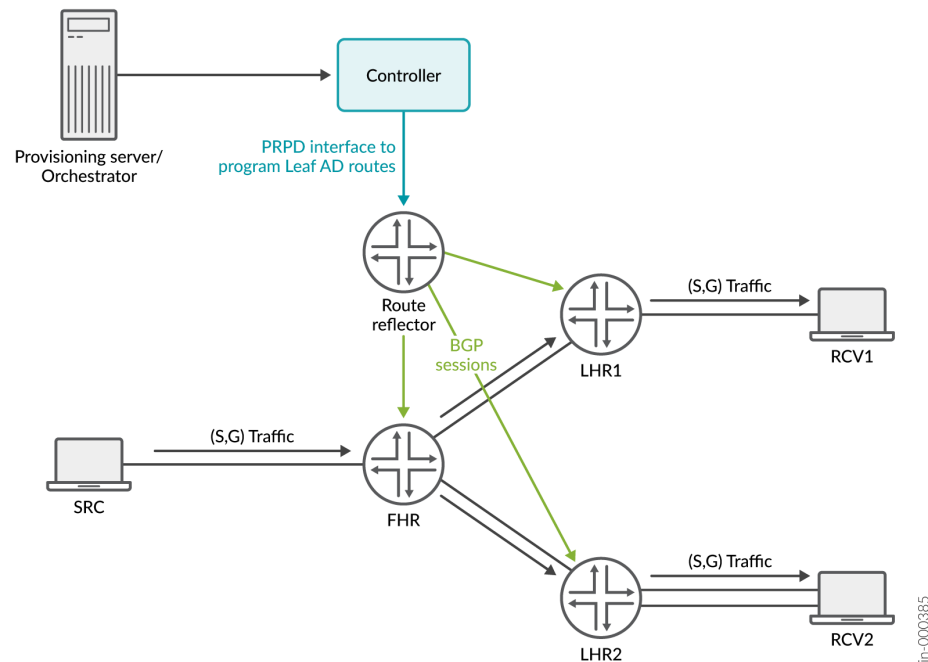
specifies the encapsulation protocols and any additional information required to use those protocols correctly.

Using this information, a router can program the multicast forwarding states to forward traffic.

Route Reflector Programming Based on PRPD and Traffic Transmission Using BGP

At times, the controller may not have PRPD connections to all the transit routers. But if BGP is running on the network, the controller needs only a single PRPD connection to a route reflector or a virtual route reflector (vRR). Using the multicast information about the source, group, and receivers, the controller programs the BGP multicast NLRI intended for all the routers on the multicast tree on the route reflector. The route reflector further propagates the BGP multicast NLRI through BGP signaling to all its neighbors which eventually reaches all the routers on the multicast tree.

Figure 150: Route Reflector Programming Based on PRPD



If the route target associated with the NLRI has the router ID of the intended router, the router accepts the route and an appropriate forwarding state is programmed on the router.



NOTE: The route target resolution RIB by default is inet.3. For BGP multicast to work, you need to explicitly set the rib resolution to inet.0 in the configuration by using this command:

```
set routing-options resolution rib bgp.rtarget.0 resolution-ribs inet.0
```

You can also apply route target constraints to filter the routes from the route reflector.

RELATED DOCUMENTATION

bgpmcast

accept-prpd-connection

inet-bgpmcast

inet6-bgpmcast

show route table

show route advertising protocol

show route receive protocol

show multicast route

7

PART

Troubleshooting

Knowledge Base | 1340

CHAPTER 29

Knowledge Base

8

PART

Configuration Statements and Operational Commands

[Junos CLI Reference Overview | 1342](#)

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- [Junos CLI Reference](#)

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- [Configuration Statements](#)
- [Operational Commands](#)