Configuration — IP Multicast Routing Protocols
Avaya Ethernet Routing Switch 8800/8600

Release 7.1
NN46205-501
Issue 05.05
September 2014

© 2014 Avaya Inc.

All Rights Reserved.

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Please note that if you acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to you by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA CHANNEL PARTNER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

License types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

“Heritage Nortel Software” means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at http://support.avaya.com/LicenseInfo/ under the link “Heritage Nortel Products”, or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants Customer a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or (in the event the applicable Documentation permits installation on non-Avaya equipment) for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those Products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the Documentation or on Avaya’s website at: http://support.avaya.com/Copyright or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

Preventing Toll Fraud

“Toll Fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya

products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and Product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for Product or Hosted Service notices and articles, or to report a problem with your Avaya Product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Chapter 1: New in this release 17
Features 17
Multicast VLAN Registration CLI enhancement 17
Chapter 2: Introduction 19
Chapter 3: IP multicast fundamentals 21
Navigation 21
Overview of IP multicast 21
SPBM MGID usage 23
Multicast host groups 24
Multicast addresses 24
Multicast protocols 25
Static source groups 25
Static IP multicast route 26
MIB support 27
Internet Group Management Protocol 29
IGMP queries 29
IGMP host reports 29
Host leave messages 30
Fast leave feature 30
Avaya Ethernet Routing Switch 8800/8600 implementation of IGMP 31
IGMP snoop 31
Multicast group trace for IGMP snoop 32
IGMP proxy 32
IGMP versions 32
IGMPv3 33
Multicast access control 36
Multicast access control policy types 37
Multicast stream limitation feature 40
Multicast Router Discovery protocol 41
Distance Vector Multicast Routing Protocol 41
Navigation 42
Reverse path forwarding 42
Pruning and grafting 42
DVMRP concepts and terminology 43
DVMRP static source groups 45
DVMRP routing policies 45
Avaya Ethernet Routing Switch 8800/8600 implementation of DVMRP 51
Protocol Independent Multicast-Sparse Mode 51
Navigation 52
PIM-SM concepts and terminology 52
Shared trees and shortest-path trees 57
Receiver joining a group 58
Receiver leaving a group 59
Source sending packets to a group 59

Required elements for PIM-SM operation ... 60
PIM-SM simplified example ... 60
PIM-SM static source groups ... 61
PIM-SMLT ... 62
PIM-SSM ... 63
SSM features ... 63
PIM-SSM architecture ... 64
Avaya Ethernet Routing Switch 8800/8600 implementation of SSM and IGMP ... 65
Avaya Ethernet Routing Switch 8800/8600 implementation of PIM-SSM over SMLT ... 68
PIM-SSM static source groups ... 69
Configuration limitations ... 69
PIM passive interfaces ... 70
PIM-SM, PIM-SSM, and DVMRP over SMLT considerations ... 71
Pragmatic General Multicast ... 71
PGM concepts and terminology ... 72
PGM network element ... 73
Multicast VLAN Registration Protocol ... 73
MVR and IGMP fast leave on Mrouter ports ... 74
Limitations to MVR ... 75
Multicast MAC filtering ... 76
High Availability ... 76
Virtualization ... 77
Limitations ... 77
Virtualization scenarios ... 77
IGMP L2 Querier ... 80
IGMP IGMP L2 Querier limitations ... 81
Chapter 4: IP multicast routing configuration ... 83
Prerequisites to IP multicast routing configuration ... 83
IP multicast routing configuration tasks ... 83
IP multicast routing configuration navigation ... 85
Chapter 5: IP multicast basic configuration using Enterprise Device Manager ... 87
IP multicast basic configuration using Enterprise Device Manager procedures ... 87
Navigation ... 88
Enabling DVMRP globally ... 89
Enabling PIM-SM globally ... 90
Enabling PIM on a brouter port ... 93
Enabling SSM globally ... 94
Enabling PIM on a VLAN interface ... 95
Enabling DVMRP on a brouter port ... 96
Enabling DVMRP on a VLAN ... 98
Configuring IGMP parameters on a brouter port ... 99
Configuring IGMP parameters on a VLAN ... 102
Chapter 6: IP multicast basic configuration using the CLI ... 107
IP multicast basic configuration procedures ... 107
IP multicast basic configuration navigation ... 109
Job aid ... 109
Configuring PIM-SM globally ... 116

Configuring PIM on a VRF ... 119
Configuring PIM on an interface ... 121
Configuring PIM on a VLAN ... 123
Configuring PIM on an Ethernet port ... 125
Configuring SSM globally ... 126
Configuring DVMRP globally ... 127
Configuring DVMRP on an interface ... 129
Configuring DVMRP on a VLAN ... 130
Configuring DVMRP on Ethernet ports ... 132
Configuring IGMP on an interface ... 133
Configuring IGMP on a VLAN ... 137
Configuring IGMP Ethernet ports ... 141
Deleting a single IP Multicast record ... 143
Chapter 7: IP multicast basic configuration using the ACLI ... 145
IP multicast basic configuration procedures ... 145
IP multicast basic configuration navigation ... 146
Job aid ... 147
Configuring PIM-SM globally ... 151
Configuring PIM on a VRF ... 155
Configuring PIM on a VLAN ... 158
Configuring PIM on an Ethernet port ... 160
Configuring SSM globally ... 161
Configuring DVMRP globally ... 162
Configuring DVMRP on a VLAN ... 165
Configuring DVMRP on Ethernet ports ... 167
Configuring IGMP on a VLAN ... 169
Configuring IGMP Ethernet ports ... 173
Deleting a single IP Multicast record ... 176
Chapter 8: DVMRP configuration using Enterprise Device Manager ... 179
Prerequisites to DVMRP configuration ... 179
Navigation ... 179
Editing DVMRP interface parameters ... 180
Editing DVMRP interface advance parameters ... 181
Viewing DVMRP neighbor parameters ... 182
Viewing DVMRP learned routes ... 183
Viewing DVMRP next-hop information ... 184
Applying the default route policy to an interface ... 185
Applying the default route policy to a VLAN ... 186
Applying the default route policy to a port ... 187
Creating a DVMRP announce policy ... 187
Applying a DVMRP announce policy to an interface ... 190
Applying a DVMRP announce policy to a VLAN ... 191
Applying a DVMRP announce policy to a port ... 192
Creating a DVMRP accept policy ... 192
Applying a DVMRP accept policy to an interface ... 194
Applying a DVMRP accept policy to a VLAN ... 194
Applying a DVMRP accept policy to a port ... 195

Applying the advertisement of local networks policy over an interface.... 195
Applying the advertisement of local networks policy over a VLAN.... 196
Applying the advertisement of local networks policy over a port.... 196
Configuring an active or passive interface type.... 197
Configuring an active or passive VLAN type.... 197
Configuring an active or passive port type.... 198
Chapter 9: DVMRP configuration using the CLI.... 201
Prerequisites to DVMRP configuration.... 201
DVMRP configuration navigation.... 201
Job aid.... 202
Showing DVMRP next hops.... 204
Applying the default route policy to an interface.... 205
Applying the default route policy to a VLAN.... 207
Applying the default route policy to a port.... 209
Creating a DVMRP policy.... 211
Applying a DVMRP announce policy to an interface.... 217
Applying a DVMRP announce policy to a VLAN.... 218
Applying a DVMRP announce policy to a port.... 219
Applying a DVMRP accept policy to an interface.... 220
Applying a DVMRP accept policy to a VLAN.... 221
Applying a DVMRP accept policy to a port.... 222
Applying the advertisement of local networks policy to an interface.... 223
Applying the advertisement of local networks policy over a VLAN.... 224
Applying the advertisement of local networks policy over a port.... 225
Creating a passive interface.... 226
Configuring an active or passive interface type.... 227
Configuring an active or passive VLAN type.... 228
Configuring an active or passive port type.... 229
Chapter 10: DVMRP configuration using the ACLI.... 231
Prerequisites to DVMRP configuration.... 231
DVMRP configuration navigation.... 231
Job aid.... 232
Showing DVMRP next hops.... 234
Applying the default route policy to a VLAN.... 235
Applying the default route policy to a port.... 237
Creating a DVMRP policy.... 239
Applying a DVMRP announce policy to a VLAN.... 246
Applying a DVMRP announce policy to a port.... 247
Applying a DVMRP accept policy to a VLAN.... 247
Applying a DVMRP accept policy to a port.... 248
Applying the advertisement of local networks policy over a VLAN.... 249
Applying the advertisement of local networks policy over a port.... 250
Configuring an active or passive VLAN type.... 251
Creating an active port.... 251
Configuring an active or passive port type.... 252
Chapter 11: PIM configuration using Enterprise Device Manager.... 253
Prerequisites to PIM configuration.... 253

Navigation ... 254
Enabling static RP ... 255
Configuring static RP ... 255
Configuring a candidate bootstrap router ... 257
Viewing the current BSR information ... 258
Configuring a PIM virtual neighbor ... 259
Changing the VLAN interface type ... 260
Editing PIM interface parameters ... 260
Viewing PIM-SM neighbor parameters ... 262
Viewing the RP set parameters ... 263
Configuring a candidate RP ... 264
Enabling square-SMLT globally ... 265
Chapter 12: PIM configuration using the CLI ... 267
Prerequisites to PIM configuration ... 267
PIM configuration navigation ... 268
Job aid ... 268
Enabling a PIM multicast border router ... 270
Configuring the PIM interface virtual neighbor ... 271
Configuring a candidate rendezvous point ... 272
Configuring static RP ... 274
Configuring a candidate BSR on an interface ... 276
Configuring a candidate BSR on an Ethernet port ... 277
Configuring a candidate BSR on a VLAN ... 278
Enabling square-SMLT globally ... 279
Chapter 13: PIM configuration using the ACLI ... 281
Prerequisites to PIM configuration ... 281
PIM configuration navigation ... 282
Job aid ... 282
Enabling a PIM multicast border router ... 283
Configuring the PIM virtual neighbor ... 284
Configuring a candidate rendezvous point ... 285
Configuring static RP ... 286
Configuring a candidate BSR on an Ethernet port ... 288
Configuring a candidate BSR on a VLAN ... 289
Enabling square-SMLT globally ... 290
Chapter 14: IGMP configuration using Enterprise Device Manager ... 291
Prerequisites to IGMP configuration ... 291
Navigation ... 291
Enabling IGMP snoop on a VLAN ... 292
Configuring IGMP interface static members ... 292
Configuring the SSM channel table ... 294
Configuring the SSM range and global parameters ... 295
Configuring multicast stream limitation on an interface ... 297
Configuring multicast stream limitation on a VLAN ... 298
Configuring multicast stream limitation on an Ethernet port ... 299
Configuring multicast stream limitation members ... 300
Adding a multicast stream limitation member ... 301

Deleting a multicast stream limitation member.... 302
Editing the IGMP interface table.... 302
Configuring IGMP sender entries.... 305
Configuring fast leave mode.... 306
Configuring multicast access control for an interface.... 308
Viewing IGMP cache information.... 309
Viewing multicast router discovery information.... 310
Viewing IGMP snoop information.... 312
Viewing IGMP group information.... 313
Chapter 15: IGMP configuration using the CLI.... 315
Prerequisites to IGMP configuration.... 315
Navigation.... 315
Job aid.... 316
Configuring multicast stream limitation on an Ethernet port.... 323
Configuring multicast stream limitation on an interface.... 325
Configuring interface multicast stream limitation members.... 327
Configuring multicast stream limitation on a VLAN.... 328
Configuring VLAN multicast stream limitation members.... 329
Configuring IGMP multicast router discovery options.... 331
Configuring IGMP multicast router discovery on a VLAN.... 332
Configuring IGMP interface static members.... 333
Configuring IGMP static members on a VLAN.... 335
Configuring SSM dynamic learning and range group.... 336
Changing the SSM range group.... 338
Configuring the SSM channel table.... 339
Configuring multicast access control for an IGMP interface.... 342
Configuring multicast access control for an IGMP Ethernet port.... 343
Configuring multicast access control for a VLAN.... 345
Configuring fast leave mode.... 346
Configuring IGMP fast leave members on a VLAN.... 347
Enabling IGMP L2 Querier globally.... 348
Configuring L2 snoop Querier address.... 349
Configuring L2 Querier on an IP interface.... 350
Configuring L2 querier message source on a interface.... 350
Viewing IP IGMP interface.... 351
Viewing IGMP snoop configuration result.... 352
Viewing IGMP configuration result.... 352
Chapter 16: IGMP configuration using the ACLI.... 355
Prerequisites to IGMP configuration.... 355
Navigation.... 355
Job aid.... 356
Configuring multicast stream limitation on an Ethernet port.... 358
Configuring multicast stream limitation on a VLAN.... 360
Configuring VLAN multicast stream limitation members.... 361
Configuring multicast router discovery options.... 362
Configuring IGMP static members.... 363
Configuring SSM dynamic learning and range group.... 364

Changing the SSM range group ... 366
Configuring the SSM channel table ... 367
Configuring multicast access control for an IGMP Ethernet port ... 369
Configuring multicast access control for a VLAN ... 370
Configuring fast leave mode ... 371
Enabling fast leave mode on a port ... 373
Configuring IGMP fast leave members on a VLAN ... 373
Enabling IGMP L2 Querier globally ... 374
Configuring L2 snoop Querier address ... 375
Resetting L2 Querier ... 375
Resetting snoop querier address ... 376
Viewing IGMP snoop configuration result ... 376
Chapter 17: PGM configuration using Enterprise Device Manager ... 379
Prerequisites to PGM configuration ... 379
Navigation ... 379
Enabling PGM globally ... 380
Configuring VLANs with PGM ... 381
Enabling PGM on an interface ... 382
Editing PGM interface parameters ... 383
Viewing PGM session parameters ... 384
Chapter 18: PGM configuration using the CLI ... 387
Prerequisites to PGM configuration ... 387
PGM configuration navigation ... 387
Job aid ... 388
Configuring PGM globally ... 389
Configuring PGM on an interface ... 390
Configuring PGM on Ethernet ports ... 392
Configuring PGM on a VLAN ... 393
Chapter 19: PGM configuration using the ACLI ... 395
Prerequisites to PGM configuration ... 395
Navigation ... 395
Job aid ... 396
Configuring PGM globally ... 396
Configuring PGM on a port or VLAN ... 398
Chapter 20: Route management using Enterprise Device Manager ... 401
Navigation ... 401
Viewing multicast route information ... 401
Viewing multicast next-hop information ... 403
Editing multicast interface information ... 404
Editing static source groups ... 405
Adding a new static source group ... 406
Configuring a static IP multicast route ... 407
Configuring IP multicast software forwarding ... 408
Configuring the mroute stream limit ... 409
Configuring the resource usage counter for multicast streams ... 410
Chapter 21: Route management using the CLI ... 413
Route management navigation ... 413

Job aid ... 413
Displaying multicast routes ... 416
Configuring a multicast route on an interface ... 416
Configuring multicast stream limits ... 417
Configuring multicast static source groups ... 419
Configuring a static IP multicast route ... 421
Configuring IP multicast software forwarding ... 422
Configuring the resource usage counter for multicast streams ... 423
Chapter 22: Route management using the ACLI ... 427
Route management navigation ... 427
Job aid ... 427
Configuring a multicast route on an interface ... 429
Configuring multicast stream limits ... 430
Configuring multicast static source groups ... 432
Configuring a static IP multicast route ... 433
Configuring IP multicast software forwarding ... 434
Configuring the resource usage counter for multicast streams ... 435
Chapter 23: Multicast flow distribution over MLT using Enterprise Device Manager .. 439
Multicast flow distribution over MLT procedures ... 439
Multicast flow distribution over MLT navigation ... 440
Configuring multicast flow distribution globally ... 440
Configuring multicast flow distribution for a multilink trunk ... 441
Chapter 24: Multicast flow distribution over MLT using the CLI ... 445
Multicast flow distribution over MLT procedures ... 445
Multicast flow distribution over MLT navigation ... 446
Job aid ... 446
Configuring multicast flow distribution globally ... 446
Configuring multicast flow distribution for a multilink trunk ... 448
Chapter 25: Multicast flow distribution over MLT using the ACLI ... 451
Multicast flow distribution over MLT procedures ... 451
Multicast flow distribution over MLT navigation ... 452
Job aid ... 452
Configuring multicast flow distribution globally ... 452
Configuring multicast flow distribution for a multilink trunk ... 454
Chapter 26: MVR configuration using Enterprise Device Manager ... 457
Navigation ... 457
Enabling MVR globally ... 457
Enabling MVR or proxy on a VLAN ... 458
Adding a receiver VLAN to the MVR VLAN ... 459
Disabling MVR ... 460
Viewing the MVR group ... 460
Viewing the MVR VLAN information ... 461
Chapter 27: MVR configuration using the CLI ... 463
Navigation ... 463
Job aid ... 463
Enabling MVR globally ... 464
Configuring MVR on a VRF ... 465

Enabling MVR or MVR Proxy on a VLAN.... 466
Adding a receiver VLAN to the MVR VLAN.... 467
Viewing MVR group information.... 468
Viewing multicast VLAN information.... 469
Chapter 28: MVR configuration using the ACLI.... 471
Prerequisites.... 471
Navigation.... 471
Job aid.... 471
Enabling MVR globally.... 472
Disabling MVR globally.... 473
Configuring MVR on VRF using ACLI.... 474
Adding a receiver VLAN to the MVR VLAN.... 474
Removing a receiver VLAN from the MVR VLAN.... 475
Enabling MVR or MVR proxy on a VLAN.... 476
Viewing MVR group information.... 477
Viewing multicast VLAN information.... 478
Chapter 29: Multicast MAC filtering using Enterprise Device Manager.... 479
Navigation.... 479
Configuring Layer 2 multicast MAC filtering.... 479
Configuring Layer 3 multicast MAC filtering.... 480
Chapter 30: Multicast MAC filtering using the CLI.... 483
Multicast MAC filtering navigation.... 483
Job aid.... 483
Configuring Layer 2 multicast MAC filtering.... 484
Configuring Layer 3 multicast MAC filtering.... 486
Chapter 31: Multicast MAC filtering using the ACLI.... 489
Multicast MAC filtering navigation.... 489
Job aid.... 489
Configuring Layer 2 multicast MAC filtering.... 489
Configuring Layer 3 multicast MAC filtering.... 491
Chapter 32: Common procedures using Enterprise Device Manager.... 493
Navigation.... 493
Configuring a prefix list.... 493
Chapter 33: Common procedures using the CLI.... 495
Job aid.... 495
Creating an IP prefix list.... 495
Chapter 34: Common procedures using the ACLI.... 499
Job aid.... 499
Creating an IP prefix list.... 499
Chapter 35: CLI show command reference.... 501
CLI show command reference navigation.... 501
DVMRP.... 502
DVMRP interface.... 503
DVMRP neighbor.... 504
DVMRP ports.... 505
DVMRP route.... 506
DVMRP VLAN.... 507

IGMP access ... 508
IGMP cache ... 509
IGMP group ... 509
IGMP interface ... 510
IGMP multicast router discovery ... 511
IGMP multicast router discovery neighbors ... 512
IGMP ports ... 512
IGMP router-alert ... 514
IGMP sender ... 514
IGMP snoop ... 515
IGMP static and blocked ports ... 515
IGMP VLAN ... 516
Layer 2 multicast MAC filters ... 517
Layer 3 multicast MAC ARP data ... 518
Multicast group trace for IGMP snoop ... 518
Multicast MLT distribution ... 519
Multicast route information ... 520
Multicast route next hop ... 522
Multicast routes on an interface ... 523
Multicast static route information ... 523
Multicast VLAN information ... 524
PGM global information ... 535
PIM active RP ... 536
PIM bootstrap router ... 537
PIM interface ... 538
PIM mode ... 539
PIM neighbor ... 540
PIM route ... 540
PIM virtual neighbor ... 542
Rendezvous points ... 542
SSM channel information ... 543
SSM group range and dynamic learning status ... 544
Static RP table ... 544
Static source groups ... 545
VLAN port data ... 546
Chapter 36: ACLI show command reference ... 547
ACLI show command reference navigation ... 547
DVMRP ... 548
DVMRP interface ... 549
DVMRP neighbor ... 550
DVMRP route ... 551
IGMP access ... 552
IGMP cache ... 553
IGMP group ... 553
IGMP interface ... 554
IGMP multicast router discovery ... 555
IGMP multicast router discovery neighbors ... 556

IGMP router-alert..... 557
IGMP sender..... 557
IGMP snoop..... 558
IGMP static and blocked ports..... 558
Layer 2 multicast MAC filters..... 559
Layer 3 multicast MAC ARP data..... 559
Multicast group trace for IGMP snoop..... 560
Multicast MLT distribution..... 561
Multicast route information..... 562
Multicast route next hop..... 563
Multicast routes on an interface..... 564
Multicast static route information..... 565
PGM global information..... 565
PIM active RP..... 566
PIM bootstrap router..... 567
PIM interface..... 568
PIM mode..... 569
PIM neighbor..... 569
PIM route..... 570
PIM virtual neighbor..... 571
Rendezvous points..... 572
SSM channel information..... 573
SSM group range and dynamic learning status..... 573
Static RP table..... 574
Static source groups..... 574
VLAN port data..... 575
Chapter 37: Customer service..... 577
Getting technical documentation..... 577
Getting Product training..... 577
Getting help from a distributor or reseller..... 577
Getting technical support from the Avaya Web site..... 578

Chapter 1: New in this release

The following sections detail what's new in *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Multicast Routing Protocols,* (NN46205-501) for Release 7.1:

• Features on page 17

Features

The following sections describe feature changes for Release 7.1.

Multicast VLAN Registration CLI enhancement

You can create and delete MVR instances on a VRF with the CLI. The default MVR instance in Global Router is disabled when you create a new MVR instance.

For information on configuring a MVR instance on VRF using CLI, see Configuring MVR on a VRF on page 465

For information on configuring a MVR instance on VRF using ACLI, see Configuring MVR on VRF using ACLI on page 474

Chapter 2: Introduction

This document describes the multicast protocols that the Avaya Ethernet Routing Switch 8800/8600 supports. For more information about the user interfaces, see *Avaya Ethernet Routing Switch 8800/8600 User Interface Fundamentals,* (NN46205-308).

For more information about troubleshooting multicast routing, see *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting ,* (NN46205-703). For more information about viewing multicast routing related statistics, see *Avaya Ethernet Routing Switch 8800/8600 Performance Management,* (NN46205-704).

Navigation

- IP multicast fundamentals on page 21
- IP multicast routing configuration on page 83
- IP multicast basic configuration using Enterprise Device Manager on page 87
- IP multicast basic configuration using the CLI on page 107
- IP multicast basic configuration using the ACLI on page 145
- DVMRP configuration using Enterprise Device Manager on page 179
- DVMRP configuration using the CLI on page 201
- DVMRP configuration using the ACLI on page 231
- PIM configuration using Enterprise Device Manager on page 253
- PIM configuration using the CLI on page 267
- PIM configuration using the ACLI on page 281
- IGMP configuration using Enterprise Device Manager on page 291
- IGMP configuration using the CLI on page 315
- IGMP configuration using the ACLI on page 355
- PGM configuration using Enterprise Device Manager on page 379
- PGM configuration using the CLI on page 387
- PGM configuration using the ACLI on page 395
- Route management using Enterprise Device Manager on page 401
- Route management using the CLI on page 413

- Route management using the ACLI on page 427
- Multicast flow distribution over MLT using Enterprise Device Manager on page 439
- Multicast flow distribution over MLT using the CLI on page 445
- Multicast flow distribution over MLT using the ACLI on page 451
- Multicast MAC filtering using Enterprise Device Manager on page 479
- Multicast MAC filtering using the CLI on page 483
- Multicast MAC filtering using the ACLI on page 489
- Common procedures using Enterprise Device Manager on page 493
- Common procedures using the CLI on page 495
- Common procedures using the ACLI on page 499
- CLI show command reference on page 501
- ACLI show command reference on page 547
- Customer service on page 577

Chapter 3: IP multicast fundamentals

IP multicast extends the benefits of Layer 2 multicasting on local area networks (LAN) to wide area networks (WAN). You can use multicasting techniques on LANs primarily to help clients and servers to find each other. With IP multicast, a source can send information to multiple destinations in a WAN with a single transmission. IP multicast results in efficiency at the source and saves a significant amount of bandwidth.

Navigation

- Overview of IP multicast on page 21
- Internet Group Management Protocol on page 29
- Multicast Router Discovery protocol on page 41
- Distance Vector Multicast Routing Protocol on page 41
- Protocol Independent Multicast-Sparse Mode on page 51
- PIM-SSM on page 63
- PIM passive interfaces on page 70
- Pragmatic General Multicast on page 71
- Multicast VLAN Registration Protocol on page 73
- Multicast MAC filtering on page 76
- High Availability on page 76
- Virtualization on page 77

Overview of IP multicast

IP multicast transmits messages to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting, except that multicasting transmits to specific groups, and broadcasting transmits to all receivers on a network. Because IP multicast transmits only one stream of data to the network where it is replicated to many receivers, multicasting saves a considerable amount of bandwidth.

IP multicast services benefit applications such as video conferencing, dissemination of datagram information, and dissemination of mail or news to a large number of recipients.

Multicast protocols use different techniques to discover delivery paths.

A distribution tree is a set of multicast routers and subnetworks that allow the members of a group to receive traffic from a source. The source of the tree depends on the algorithm used by the multicast protocol. The following figure is an example of a simple distribution tree where S is the multicast source, and the arrows indicate the multicast broadcast procedure.

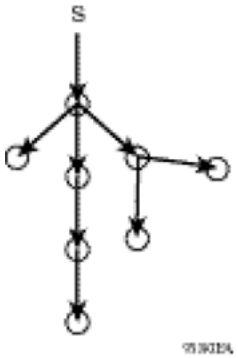

Figure 1: Multicast distribution tree and broadcasting

Broadcast and prune methods use multicast traffic to build the distribution tree. Periodically, data is sent out or broadcast from the source to the extremities of the internetwork to search for active group members. If no local members of the group exist, the router sends a message to the host, removing itself from the distribution tree, and thus pruning the router.

The following figure illustrates how routers are pruned from the distribution tree. First, a message is sent to the source, after which the pruned routers do not receive multicast data.

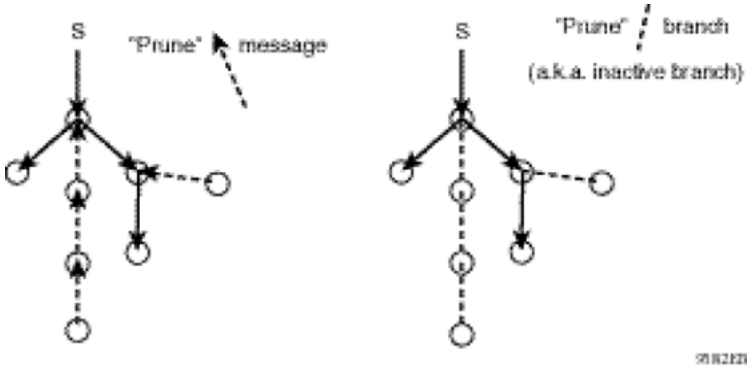

Figure 2: Pruning routers from a distribution tree

Reverse path multicast is based on the concept that a multicast distribution tree is built on the shortest path from the source to each subnetwork containing active receivers. After a datagram arrives on an interface, the router determines the reverse path to the source of the datagram by examining the routing table of known network sources. If the datagram is not on the optimal delivery tree, it is discarded.

Multicast host groups and their group members permit the IP multicast router to transmit just to those groups interested in receiving the traffic. The Avaya Ethernet Routing Switch 8800/8600 uses the Internet Group Management Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. A router communicates with the hosts on a local network by sending IGMP queries. Hosts respond by issuing IGMP reports. For more

information about host groups, see Multicast host groups on page 24 and Multicast addresses on page 24. For more information about IGMP, see Internet Group Management Protocol on page 29.

Multicast traffic forwarding transmits frames to all interfaces or subnets for which IGMP reports are received for the multicast group indicated in the destination IP address. Multicast packets forwarded within the same virtual LAN (VLAN) remain unchanged. Packets are not forwarded to networks that do not use members of the multicast group indicated in the destination IP address.

Important:

Multicast traffic forwarding on the Avaya Ethernet Routing Switch 8800/8600 does not forward jumbo frames.

SPBM MGID usage

The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a specific port number, the data is directed to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs. The system also reserves a small number of MGIDs.

SPBM also requires MGIDs for proper operation. When SPBM is enabled on the switch, the system reserves 519 MGIDs for SPBM operation. Therefore, the number of MGIDs on the system available for VLANs and IP multicast traffic is reduced by 519. To determine how many MGIDs are available, enter **show sys mgid-usage**.

```
ERS-8606:5# show sys mgid-usage
*******************************************************************************
Command Execution Time: WED FEB 02 19:42:27 2011 UTC
*******************************************************************************
Number of MGIDs used for VLANs : (3)
Number of MGIDs used for SPBM : (519)
Number of MGIDs used for multicast : (0)
Number of MGIDs remaining for VLANs : (1514)
Number of MGIDs remaining for multicast : (2048)
```

Before you enable SPBM on the switch, be sure that your network will not be adversely affected by this reduction in available MGIDs.

The Ethernet Routing Switch 8800/8600 supports a total of 4096 MGIDs, split between the system, VLAN, IPMC, and now SPBM. You can reserve MGIDs for IP Multicast (IPMC) traffic. You can reserve between 64 and 4084 MGIDs for IPMC. The default for IPMC is 2048. It is the responsibility of the network administrator to fully understand the network deployment strategy. Please ensure that MGIDs are planned appropriately. If assistance is required, please contact your Avaya technical representative.

For information about reserving MGIDs for IPMC, see *Avaya Ethernet Routing Switch 8800/8600 Administration* (NN46205–605).

Multicast host groups

IP multicast is a method for addressing, routing, and delivering a datagram to a collection of receivers called a host group.

Host groups are permanent or transient, with the following characteristics:

- A permanent host group uses a well-known, administratively assigned IP multicast group address. This address is permanent and defines the group. A permanent host group can consist of zero or more members.
- A transient host group exists only as long as members need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

A host system on an IP network sends a message to a multicast group by using the IP multicast address for the group. To receive a message addressed to a multicast group, however, the host must be a member of the group and must reside on a network where that group is registered with a local multicast router.

An IP multicast host group can consist of zero or more members and places no restrictions on its membership. Host members can reside anywhere, they can join and leave the group, and they can be members of more than one group at the same time.

In general, hosts that are members of the same group reside on different networks. However, a range of multicast addresses (224.0.0.x) is reserved for groups that are locally scoped. All message traffic for these hosts typically remains on the local network. Hosts that belong to a group in this address range and that reside in different networks do not receive message traffic for each other.

Important:

With the Avaya Ethernet Routing Switch 8800/8600, you can apply a special set of filters (global filters) to multicast packets. You can create, deny, or accept filters to configure the sources that can receive and send data.

Multicast addresses

Each host group uses a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP version 4 multicast address is a Class D address (the high-order bits are set to 1110) from 224.0.1.0 to 239.255.255.255. These addresses are assigned statically for use by permanent groups and dynamically for use by transient groups.

Routing protocols and other low-level protocols reserve the block of addresses from 224.0.0.1 to 224.0.0.255. Multicast routers do not forward datagrams with addresses in this range because the time-to-live (TTL) value for the packet is usually set to 1.

Multicast protocols

You can use the following protocols to enable multicasting on an Avaya Ethernet Routing Switch 8800/8600:

- Internet Group Management Protocol (IGMP)—learns the existence of host group members on directly attached subnets
- Multicast Router Discovery (MRDISC) protocol—discovers multicast routers in a Layer 2 bridged domain configured for IGMP snoop
- Distance Vector Multicast Routing Protocol (DVMRP)—a dense-mode protocol suitable for implementation in networks that are densely populated by receivers
- Protocol Independent Multicast (PIM)
  - Sparse Mode (PIM-SM) protocol—suitable for implementation on networks that are sparsely populated by receivers
  - Source Specific Multicast (PIM-SSM) protocol—uses a one-to-many model where members can only receive traffic from a single source

    This protocol is suitable for television channels and other content-distribution applications
- Pragmatic General Multicast (PGM) Protocol—suitable for multicast applications that require reliable, ordered, duplicate-free delivery of multicast traffic

Static source groups

Use static source groups to configure static source-group entries in the DVMRP, PIM-SM, or PIM-SSM multicast routing table. DVMRP and PIM cannot prune these entries from the distribution tree. Even if no receivers for the group exist, the multicast stream for a static source-group entry stays active. DVMRP and PIM never prune static forwarding entries. When you no longer need the entries, you manually delete them.

To configure static source groups, you must first globally enable either DVMRP or PIM. If you disable DVMRP or PIM, the switch saves all of the configured static source-group entries and deactivates them. After you reenable DVMRP or PIM, the switch reactivates the static source groups.

Static source groups ensure that the multicast route (mroute) records remain in the distribution tree. When receivers join the group, they do not experience a delay in receiving multicast data because they do not need to graft onto the group or start a join process in the case of PIM. This timing is essential for applications that must send the multicast data to a receiver as soon

as the receiver joins the group, for example, when a switch delivers television channels to receivers. After the receiver turns the channel, which is equivalent to joining a group, the receiver can view the channel immediately.

Static entries result in continuous traffic if the source is active, even when no receivers are present. However, traffic is not forwarded by an Avaya Ethernet Routing Switch 8800/8600 with a static entry if no receivers exist, but is forwarded continuously to the switch where the entry is programmed and crosses intermediate switches on the path.

You can configure static source-group entries for a specific source or subnet. If several sources on the same subnet send traffic to the same group, traffic for all these sources flows continuously when you use the subnet configuration.

After you configure static source groups, be aware of the following switch actions:

- If you disable DVMRP or PIM, the switch deactivates all of the static source groups. After you reenable DVMRP or PIM, the switch reactivates the static source groups.
- In DVMRP and PIM-SM configurations, the static source-group feature works for both specific source addresses and subnet addresses by using the SrcSubnetMask field.

  If the network mask is 255.255.255.255, the full source address is used to match the (S,G), which is the specific source case. When you configure the network mask field as a subnet mask for the source, only the source subnet is used to match (S,G)s.
- In PIM-SSM configurations, static source groups use the following limitations:

  - Subnets: SSM static source groups work only with specific IP addresses. This means that static source groups cannot work with source subnets, so the mask must use a full 32-bit mask, 255.255.255.255, and the source must use a host address.
  - SSM channels: Static source groups cannot conflict with SSM channels and vice versa. After you configure a static source group or an SSM channel, the switch performs a consistency check to make sure no conflicts exist. You cannot map one group (G) to different sources, for both a static source group and an SSM channel.

    If a group is already mapped to a source and you try to map it to a different source, the switch detects the conflict and an error message appears. For example, if G1 is already defined in the SSM channel table as (S1,G1), you cannot configure G1 as static source group (S2,G1). However, you can configure the same entry (S1,G1) both in the SSM channel table and as a static source group. If no conflict exists between the two tables, the configuration is allowed.

Static IP multicast route

Release 5.1 introduces a new static IP multicast route (mroute) to separate the paths for unicast and multicast streams. Only PIM-SM and PIM-SSM use this mroute. Adding an mroute does not affect the switching or routing of unicast packets.

The entries in the mroute use the following attributes:

- IP prefix or IP mask—the destination network for the added route
- Reverse Path Forwarding (RPF) address—the IP address of the RPF neighbor towards the rendezvous point (RP) or source
- route preference—the administrative distance for the route

  If the unicast routing table and the multicast-static IP route use different routes for the same destination network, the system compares the administrative distance with that of the protocol that contributed the route in the unicast routing table.
- route status—the status, either enabled or disabled, of the route

The system does not advertise or redistribute multicast-static IP routes. The system uses these routes only for RPF calculation. The system uses the following rules to determine RPF:

- Direct or local routes for a destination take precedence over a route for the same destination in the static route table.
- If a route exists in the static route table, and no route exists in the unicast routing table for the destination, the system uses the route in the static route table.
- If a route is available in both the unicast routing table and the static route table, the system uses the route from the static route table only if the administrative distance is less than or equal to that of the unicast route entry.
- If no route exists in the static route table for the destination, the system uses the route from the unicast routing table, if available.
- The system performs a longest prefix match during a lookup in the static route table. The lookup ignores routes that are administratively disabled.
- After the system performs a lookup within the static mroute table, if multiple routes exist for a matching prefix, the system chooses the route with the least preference. If multiple routes exist with a matching prefix and the same preference, the system chooses the route with the highest RPF address. This selection method only occurs within the static mroute table; the system still compares the selected route with a route from RTM, if one exists.

Because the switch uses two types of static IP route, one for unicast routes and one for multicast routes, the CLI command **`show ip route info`** does not provide the correct RPF information to a source or RP. Avaya recommends that you use the **`show ip mroute rpf`** command.

MIB support

Release 5.1 standardizes the mandatory objects in the multicast management information bases (MIB). The following table identifies the tables and groups supported for each standard MIB. For more information about each MIB, see the RFC.

Table 1: Standard multicast MIB support

Supported tables and groups	Exceptions	Notes
RFC 2932—IPv4 Multicast Routing MIB		
ipMRouteTable	ipMRoutePkts, ipMRouteDifferentInIfPackets, ipMRouteOctets	—
IP Multicast Routing MIB-Group	—	The value of ipMRouteEnable always appears as enabled.
ipMRouteNextHopTable	—	—
ipMRouteInterfaceTable	—	The system does not perform an interface-specific rate limit. The value always appears as 0.
RFC 2933—Internet Group Management Protocol MIB		
igmpInterfaceTable	igmpInterfaceProxyIfIndex	—
igmpCacheTable	—	A set operation for the igmpCacheSelf object results in failure. The value of this object is always false. The router never becomes a member of a forward-able multicast group address. This cache table does not store protocol specific link-local multicast addresses. Avaya supports read-only access to igmpCacheStatus. The IGMP cache table displays only dynamically learned entries; you cannot create or delete the table or a row within the table.
RFC 2934—Protocol Independent Multicast MIB for IPv4		
pimInterfaceTable	—	—
pimNeighborTable	—	—
pimRPSetTable	—	—
pimCandidateRPTable	—	—
pimComponentTable	—	—

Supported tables and groups	Exceptions	Notes
pimIpMRouteTable	—	Avaya support includes only the following: pimIpMRouteUpstreamAssertTimer and pimIpMRouteFlags

Internet Group Management Protocol

The IGMP has the following characteristics:

- A host uses IGMP to register group memberships with the local querier router to receive datagrams sent to this router that are targeted to a group with a specific IP multicast address.
- A router uses IGMP to learn the existence of group members on networks to which it is directly attached. The router periodically sends a general query message to each of its local networks. A host that is a member of a multicasting group identifies itself by sending a response.

IGMP queries

When multiple IGMP routers operate on a network, one router is elected to send queries. This elected querier periodically sends host membership queries (also known as general queries) to its attached local subnets. The Avaya Ethernet Routing Switch 8800/8600 supports queries from all three versions of IGMP.

IGMP host reports

A host that receives a membership query from a local router can respond with a host membership report, one for each multicast group that joins. A host that receives a query delays its reply by a random interval and listens for a reply from other hosts in the same host group. For example, consider a network that includes two host members—host A and host B—of the same multicast group. The router sends out a host membership query on the local network. Both host A and host B receive the query and listen on the network for a host membership report. The delay timer for host B expires first, so it responds to the query with a membership report. Hearing the response, host A does not send a report of its own for the same group.

Each query from a router to a host includes a Maximum Response Time field. IGMP inserts a value n into this field specifying the maximum time in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and n tenths of a second for the period that it waits before sending a response. This calculation is

true for IGMP version 2 and 3. For IGMP version 1, this field is set to 0 but defaults to a value of 100, that is, 10 seconds.

If at least one host on the local network specifies that it is a member of a group, the router forwards to that network all datagrams bearing the multicast address for the group.

Upon initialization, the host can immediately issue a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same as requested reports.

After hosts and routers are in a steady state, they communicate in a way that minimizes the exchange of queries and reports. The designated routers set up a path between the IP multicast stream source and the end stations and periodically query the end stations about whether to continue participation. As long as a client continues to participate, all clients, including nonparticipating end stations on the switch port, receive the IP multicast stream.

Host leave messages

When an IGMPv2 host leaves a group and it is the host that issues the most recent report, it also issues a leave group message. The multicast router on the network issues a group-specific query to determine whether other group members exist on the network. If no host responds to the query, the router assumes that no members belonging to that group exist on that interface.

Fast leave feature

The Avaya Ethernet Routing Switch 8800/8600 supports a fast leave feature that is useful for multicast-based television distribution applications. Fast leave relies on an alternative leave process where the switch stops sending traffic for the group immediately after receiving a leave message, without issuing a query to check if other group members are present on the network. Fast leave alleviates the network from additional bandwidth demand when you change television channels.

Fast leave mode

The Avaya Ethernet Routing Switch 8800/8600 provides several fast leave processes for IP multicast:

- immediate leave with one user for each interface
- immediate leave with several users for each interface
- standard IGMP leave based on a Last Member Query Interval (LMQI), which is configurable in tenths of seconds

Fast leave modifies the IGMP leave processing mechanism on an IGMP interface. After receiving an IGMP leave on a fast leave enabled interface, the switch does not send a group-

specific query and immediately stops sending traffic to the leaving member (IGMP host) port. Without fast leave, traffic is forwarded until the group-specific query times out. This wastes bandwidth if no receiver interested in the group traffic exists.

Fast leave mode provides two options of the fast leave mechanism—single-user mode and multiple-users mode.

- Single-user mode: In this mode, the port stops receiving traffic immediately after a group member on that port sends a leave message. Avaya recommends that you use the single-user mode when each switch interface port connects to only one IGMP host.
- Multiple-users mode: Use this mode if the switch interface port is connected to multiple IGMP hosts. In this case, the port stops receiving traffic after all members leave the IGMP group. The switch removes the leaving IGMP member and, if more group members exist on that port, the switch continues sending traffic to the port.

  When it runs in multiple-users mode, the Avaya Ethernet Routing Switch 8800/8600 must use the correct membership information. To support multiple-users mode, multicast receivers on the same interface cannot use IGMP report suppression. If you must use IGMP report suppression, Avaya recommends that you do not use this mode. Instead, use the LMQI (configurable in units of 1/10ths of seconds) to provide a faster leave process while still sending group-specific queries after a leave message is received.

Fast leave mode applies to all fast leave enabled IGMP interfaces.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces. Although Internet Group membership Authentication Protocol (IGAP) interfaces are always fast leave enabled, they ignore this mode because they only operate in the multiple-users mode.

Avaya Ethernet Routing Switch 8800/8600 implementation of IGMP

You can enable and disable multicast routing on an interface basis. If you disable multicast routing on an interface, IGMP queries are not generated. If the switch or interface is in IGMP router behavior mode, for example, DVMRP or PIM enabled, you cannot configure IGMP snoop. The switch still learns the group membership and snoops multicast receivers on the switch VLAN or ports.

IGMP snoop

The Avaya Ethernet Routing Switch 8800/8600 provides IP multicast capability when you use it as a switch. Functioning as a switch, it supports all three versions of IGMP to prune group membership for each port within a VLAN. This feature is called IGMP snoop.

Important:

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

Use the IGMP snoop feature to optimize the multicast data flow, for a group within a VLAN, to only those ports that are members of the group. The switch builds a database of group members by listening to IGMP reports from each port. It suppresses the reports heard by not forwarding them to ports other than the one receiving the report, thus forcing the members to continuously send their own reports. The switch relays group membership from the hosts to the multicast routers. It forwards queries from multicast routers to all port members of the VLAN. Furthermore, the switch forwards multicast data only to the participating group members and to the multicast routers within the VLAN.

Multicast group trace for IGMP snoop

This feature helps you to monitor the multicast group trace for an IGMP snoop-enabled Avaya Ethernet Routing Switch 8800/8600. You can view the multicast group trace from the CLI or Enterprise Device Manager.

Multicast group trace tracks the data flow path of the multicast streams. Group trace tracks information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port.

IGMP proxy

If an Avaya Ethernet Routing Switch 8800/8600 receives multiple reports for the same multicast group, it does not transmit each report to the multicast upstream router. Instead, the switch consolidates the reports into a single report and forwards it. If the switch adds another multicast group or receives a query since the last upstream report transmission, the switch forwards the report onto the multicast router ports. This feature is called IGMP proxy.

IGMP versions

The Avaya Ethernet Routing Switch 8800/8600 supports IGMPv1, IGMPv2, and IGMPv3. IGMPv1 and IGMPv2 are backward compatible and can exist together on a multicast network.

Beginning with Release 5.1, IGMPv3 for PIM-SSM is backward compatible with IGMPv2. The following list describes the purpose for each version:

- IGMPv1 provides the support for IP multicast routing. IGMPv1 specifies the mechanism for communicating IP multicast group membership requests from a host to its locally attached routers. For more information, see RFC 1112.
- IGMPv2 extends the features in IGMPv1 by quickly reporting group membership termination to the routing protocol. This feature is important for multicast groups with highly volatile group membership. For more information, see RFC 2236.
- IGMPv3 supports the PIM Source Specific Multicast (SSM) protocol. A host can selectively request or filter traffic from individual sources within a multicast group. For more information, see RFC 3376.

IGMPv3

For the Avaya Ethernet Routing Switch 8800/8600 implementation of PIM-SSM, only one source is supported per group.

The switch accepts IGMPv3 reports with the following restraints:

- The group record is a current-state record or source-list-change record. Filter-mode-change record is not implemented.
- A current-state record can only use a record type of MODE_IS_INCLUDE. A group record in the current-state record can use no more than one source. MODE_IS_EXCLUDE is not implemented.
- A source-list-change record can only use a record type of BLOCK_OLD_SOURCES. The record can use no more than one source. ALLOW_NEW_SOURCES is not implemented.
- If a group already uses a source associated with it through either static configuration or dynamic learning, a current-state record with a different source is discarded with a message and log.
- The equivalent for the leave function in IGMPv2 is achieved by a source-list-change record with a record type of BLOCK_OLD_SOURCES.
- After the interface receives an IGMP query with a lower version, the IGMP interface operates at the lower version with a message and log. After the interface receives an IGMP report with a lower version, the interface discards the report with a message and log.
- You can enable IGMPv3 only after you enable SSM on the interface. If you do not enable SSM, you cannot enable IGMPv3.
- After you enable dynamic learning for the SSM source-group table, the table learns source-group pairs from IGMPv3 reports. The learned entry times out after the corresponding report times out due to IGMP leave or timer expiration.

The static member and access list is indexed by group address only because the group address to channel uses a one-to-one match. You cannot configure a static member on an SSM group that does not exist in the SSM source group table.

After you enable IGMPv3, the following actions occur:

- The switch does not accept IGMP packets with a group address out of the SSM range.
- The switch associates each group with only one source as indicated in the SSM source group table.
- The switch processes only reports with type MODE_IS_INCLUDE and BLOCK_OLD_SOURCE. The switch discards reports with other types.
- The switch discards IGMPv1 or IGMPv2 reports that it receives.
- Existing IGMP features such as static member, access list, and pseudo report are based on group rather than channels.
- You can enable IGMPv3 only after PIM-SSM or SSM-snoop is enabled.
- After you change the version on an interface to or from IGMPv3, you disrupt existing multicast traffic on that interface. Avaya recommends that you do not make this change when the switch runs multicast traffic.

IGMPv3 backward compatibility

Beginning with Release 5.1, IGMPv3 for PIM-SSM is backward compatible with IGMPv2. In previous releases, if you configure the Avaya Ethernet Routing Switch 8800/8600 to use IGMPv3, the switch supports IGMPv2 and IGMPv1 query messages only; the switch discards other messages types.

In Release 5.1 and later, you can configure the switch to operate in v3-only mode, in v2-v3 compatibility mode, or both modes at once. If you configure the switch to use only v3-only mode, it ignores all v2 and v1 messages except the query message, as in previous software releases.

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv1, v2, and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message; if it is a v2 message, PIM-SM or IGMP snoop processes handle the message.

After the switch receives an IGMPv2 leave message, and the group address in it is within SSM range, the switch sends the group-and-source specific query. If the group address is not within the SSM range, the switch sends the group specific query.

According to RFC 3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.
- If a router that is not explicitly configured to use IGMPv1 or IGMPv2, hears an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure if the Avaya Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only

capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

In v2-v3 compatibility mode, an IGMPv2 host can only join the static group in SSM range when the interface operates in PIM-SSM mode or IGMP SSM mode as a Layer 2 querier.

IGMPv3 backward compatibility supports partial High Availability. You can use the compatibility mode with Split MultiLink Trunking (SMLT). One core switch sends an SMLT message to the other core switch after it receives an IGMPv3 message. This action synchronizes the IGMP host information.

Multicast virtualization support includes IGMPv3 backward compatibility.

The following two paragraphs discuss the interoperability of the IGMPv3 backward compatibility feature with Ethernet Switches and Ethernet Routing Switches that do not support IGMPv3.

When you enable snoop on Ethernet Switches and Ethernet Routing Switches that do not support IGMPv3, they do not process v3 queries. The Ethernet Routing Switch 8800/8600 sends only IGMPv3 queries if you configure IGMPv3, even if you enable backward compatibility. If an Ethernet Switch or Ethernet Routing Switch with snoop enabled exists on a Layer 2 VLAN between the IGMPv2 client and the Ethernet Routing Switch 8800/8600, the host membership reports that the IGMPv2 client sends do not reach the Ethernet Routing Switch 8800/8600 unless you configure static router ports in the IGMP configuration for the VLAN on the switch.

You may not immediately notice this issue if you change the Avaya Ethernet Routing Switch 8800/8600 from IGMPv2 to IGMPv3. After this change occurs, the Ethernet Switch or Ethernet Routing Switch still sends the membership reports towards the Ethernet Routing Switch 8800/8600 for a period of time because it still knows where the router is based off the last v2 query. Eventually, the Ethernet Switch or Ethernet Routing Switch stops sending host membership reports because it only receives and drops IGMPv3 queries from the Ethernet Routing Switch 8800/8600.

A limitation of the IGMP dynamic downgrade feature on the Ethernet Routing Switch 8800/8600 exists with Avaya Ethernet Switches and Ethernet Routing Switches when you use them only as Layer 2 devices because these switches only send queries on the management VLAN. The Ethernet Routing Switch 8800/8600 can only dynamically downgrade the IGMP version for a VLAN if it receives a query on that particular VLAN. If multiple VLANs exist between the Ethernet Routing Switch 8800/8600 and one of these switches, dynamic downgrade only works on the management VLAN.

IGMPv3 downgrades

Do not change the IGMP version at the interface level from IGMPv3 to IGMPv2 and back to IGMPv3, as this is not supported. When the IGMP version is downgraded from v3 to v2, this sets the operating version of all switches in the network to IGMPv2, and all IGMPv3 membership reports are discarded. This also deletes the IGMP group table information. This

is in accordance with the RFC, which states that downgrading is supported but upgrading is not.

Multicast access control

Multicast access control is a set of features unique to the Avaya Ethernet Routing Switch 8800/8600 that operate with standard existing multicast protocols. You can configure multicast access control on an IP multicast-enabled port or VLAN with an access control policy that consists of several IP multicast groups.

You can use this feature to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams). For example, in a television distribution application, instead of applying a filter to each channel (multicast group), you can apply a multicast access policy to a range of channels (groups), thereby reducing the total number of filters and providing a more efficient and scalable configuration. Also, if you want to add or remove television channels from a package, you can modify the multicast access policy; you do not need to change filters for individual VLANs or ports. Multicast access policies contain an ID and a name (for example, PremiumChannels), the list of IP multicast addresses, and the subnet mask.

Multicast access control is not a regular filtering configuration. Multicast access control is specifically designed for multicast streams and relies on handling multicast control and initial data to prevent hosts from sending or receiving specified multicast streams; it does not use filters. Also, multicast access control provides a list of multicast groups in one configuration using the same routing policy prefix list configuration. For more information about prefix lists, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing,* (NN46205-523). You can configure multicast access control, and change it dynamically to support changes in the configuration without restarting the protocol. You can change the access capabilities of a user or service subscriber without loss of service.

The following paragraph describes a typical application.

Your local cable television company offers three packages; each one includes 35 channels (35 multicast groups). Each package is configured in an access control policy. This policy is applied to a set of VLANs or ports to prevent users from viewing the channels on those VLANs. Use the same policy to prevent users from sending traffic to those groups (also known as spoofing) by specifying the deny-tx option for that port. After you define the packages, you can use them for access policy configuration. You can easily change the package by changing the group range, without changing all the port configurations.

The multicast access control functionality applies to an IP multicast application where you require user access control. You can use it in financial-type applications and other enterprise applications, such as multicast-based video conferencing.

Multicast access control policy types

Six types of multicast access control policies exist:

- deny-tx
- deny-rx
- deny-both
- allow-only-tx
- allow-only rx
- allow-only-both

The tx policies control the sender and ingress interface for a group; the rx policies control the receivers and egress interface for a group.

deny-tx

Use the deny-tx access policy to prevent a matching source from sending multicast traffic to the matching group on the interface where the deny-tx access policy is configured. Configure this policy on the ingress interface to the multicast source. The deny-tx access policy performs the opposite function of the allow-only-tx access policy. Therefore, the deny-tx access policy and the allow-only-tx access policy cannot exist on the same interface at the same time.

For example, in Figure 3: Data flow using deny-tx policy on page 37, a deny-tx access policy is configured on VLAN 1 (the ingress VLAN to the Avaya Ethernet Routing Switch 8800/8600). This policy prevents multicast traffic sent by Sender from forwarding from VLAN 1 to a receiver, consequently preventing Receiver 1 and Receiver 2 from receiving data from the multicast group. You can create receive-only VLANs, such as VLAN 1, with the deny-tx policy.

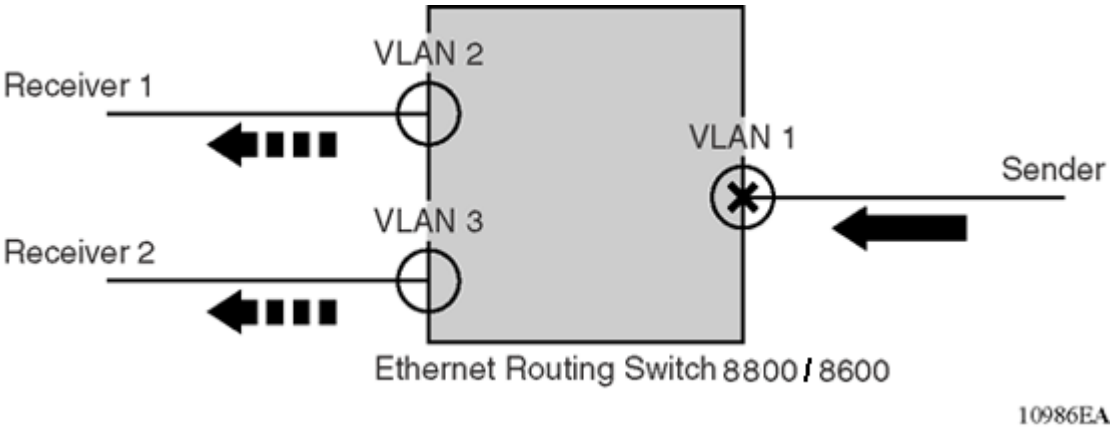

Figure 3: Data flow using deny-tx policy

deny-rx

Use the deny-rx access policy to prevent a matching group from receiving IGMP reports from the matching receiver on the interface where the deny-rx access policy is configured. The deny-rx access policy performs the opposite function of the allow-only-rx access policy. Therefore, the deny-rx access policy and the allow-only-rx access policy cannot exist on the same interface at the same time.

For example, in Figure 4: Data flow using deny-rx policy on page 38, a deny-rx access policy is configured on VLAN 2, preventing IGMP reports sent by Receiver 1 from receiving on VLAN 2. You can deny a multicast group access to a specific VLAN or receiver using the deny-rx policy.

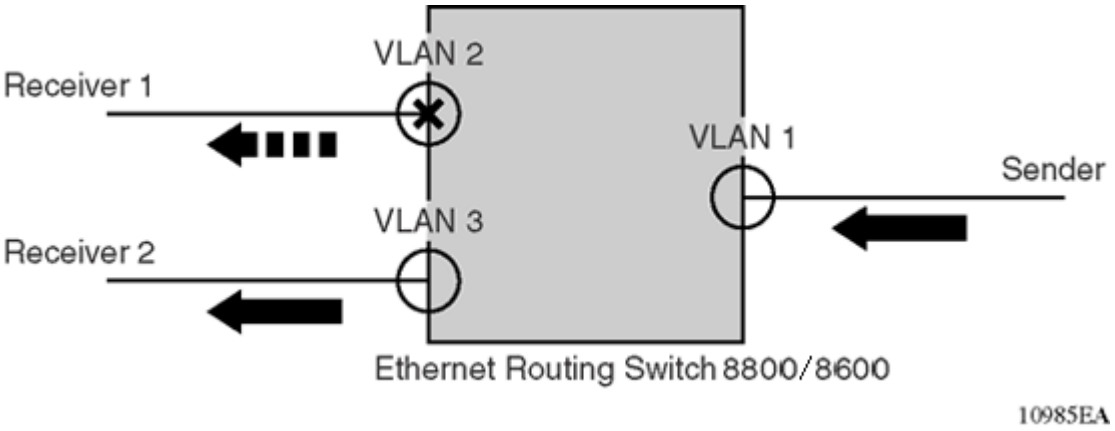

Figure 4: Data flow using deny-rx policy

deny-both

Use the deny-both access policy to prevent a matching IP address from both sending multicast traffic to, and receiving IGMP reports from, a matching receiver on an interface where the deny-both policy is configured. You can use this policy to eliminate all multicast activity for a receiver or source in a specific multicast group. The deny-both access policy performs the opposite function of the allow-only-both access policy. Therefore, the deny-both access policy and the allow-only-both access policy cannot exist on the same interface at the same time.

For example, in Figure 5: Data flow using deny-both policy on page 39, a deny-both access policy is configured on VLAN 2, preventing VLAN 2 from receiving IGMP reports sent by Receiver 2, and preventing multicast traffic sent by Sender 2 from forwarding from VLAN 2. You can prevent certain VLANs from participating in an activity involving the specified multicast groups with the deny-both policy.

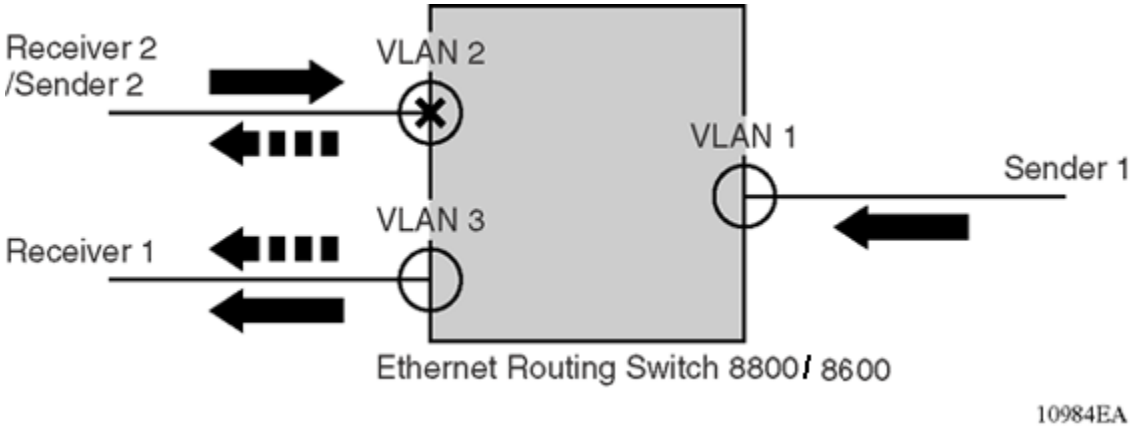

Figure 5: Data flow using deny-both policy

allow-only-tx

Use the allow-only-tx policy to allow only the matching source to send multicast traffic to the matching group on the interface where the allow-only-tx policy is configured. The switch discards all other multicast data received on this interface. The allow-only-tx access policy performs the opposite function of the deny-tx access policy. Therefore, the allow-only-tx access policy and the deny-tx access policy cannot exist on the same interface at the same time.

allow-only-rx

Use the allow-only-rx policy to allow only the matching group to receive IGMP reports from the matching receiver on the interface where the allow-only-rx access policy is configured. The switch discards all other multicast data received on this interface. The allow-only-rx access policy performs the opposite function of the deny-rx access policy. Therefore, the allow-only-rx access policy and the deny-rx access policy cannot exist on the same interface at the same time.

allow-only-both

Use the allow-only-both policy to allow only the matching IP address to both send multicast traffic to, and receive IGMP reports from, the matching receiver on the interface where the allow-only-both access policy is configured. The switch discards all other multicast data and IGMP reports received on this interface. The allow-only-both access policy performs the opposite function of the deny-both access policy. Therefore, the allow-only-both access policy and the deny-both access policy cannot exist on the same interface at the same time.

Host addresses and masks

When you configure multicast access policies, you must specify the host (IP) address and host (subnet) mask of the host that is filtered (the host that sends multicast traffic).

You can use the host subnet mask to restrict access to a portion of the host network. For example, if you configure the host subnet mask as 255.255.255.255, you use the full host address. To restrict access to a portion of the network of a host, use a subnet mask such as 255.255.255.0. Access control applies to the specified subnet only.

Multicast stream limitation feature

You can configure the multicast stream limitation feature to limit the number of multicast groups that can join a VLAN. By limiting the number of concurrent multicast streams, a service provider can, for example, protect the bandwidth on a specific interface and control access to multicast streams.

Use multicast stream limitation in an environment where you want to limit users to a certain number of multicast streams simultaneously. For example, a television service provider can limit the number of television channels a user can watch at a time. (To a television service provider, a multicast stream is synonymous with a television channel.) If a user purchases a service contract for two television sets, they can use two channels flowing at the same time, but not a third. The service provider can control the bandwidth usage in addition to preventing users from watching more than the allowed number of channels at a point in time.

You can enable the multicast stream limitation feature on the Avaya Ethernet Routing Switch 8800/8600 by using one of the following methods:

- for each interface: This limitation controls the total number of streams for all clients on this interface.
- for each interface port: This limitation controls the number of streams for all clients on this interface port.
- for each Ethernet port: This limitation controls the number of streams for all clients on this Ethernet port.
- for each VLAN: This limitation controls the total number of streams for all clients on this VLAN. This method is equivalent to the interface stream limitation.
- for each VLAN port: This limitation controls the number of streams for all clients on this VLAN port. This method is equivalent to the interface port stream limitation.

You can configure the maximum number of streams for each limit independently. After the number of streams equals the limit, the interface, port, or VLAN drops additional join reports for new streams.

Multicast Router Discovery protocol

The Multicast Router Discovery (MRDISC) protocol automatically discovers multicast-capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and IGMP host membership reports. This feature is useful in a Layer 2 bridging domain that is configured for IGMP snoop.

IGMP multicast router discovery consists of three message types that discover multicast routers on the network:

- Multicast router advertisements sent by routers to advertise that IP multicast forwarding is enabled on an interface.
- Multicast router solicitations sent by routers to solicit a response of multicast router advertisements from all multicast routers on a subnet.
- Multicast router termination messages sent when a router terminates its multicast routing functions.

Multicast routers send multicast router advertisements periodically on all interfaces where multicast forwarding is enabled. Multicast routers also send advertisements in response to multicast router solicitations.

Multicast routers send solicitations to the IGMP-MRDISC all-routers multicast group that uses a multicast address of 224.0.0.2. Multicast routers send solicitations whenever a router needs to discover multicast routers on a directly attached subnet.

Multicast routers send termination messages after a router terminates its multicast routing functions. Other non-IP forwarding devices, such as Layer 2 switches, can send multicast router solicitations to solicit multicast router advertisements.

After you enable IGMP snoop on an Avaya Ethernet Routing Switch 8800/8600, MRDISC is enabled by default.

Important:

The Multicast Router Discovery protocol is not supported on brouter ports.

Distance Vector Multicast Routing Protocol

The Distance Vector Multicast Routing Protocol (DVMRP) is a distance vector type of multicast routing protocol. It advertises shortest-path routes to multicasting source networks, that is, a network containing hosts that can issue multicast datagrams. In this respect, DVMRP is the opposite of the Routing Information Protocol (RIP), which advertises all routes to destination

networks. Coupled with IGMP, DVMRP learns membership for a multicast stream from both the routers and directly attached hosts.

DVMRP constructs a different distribution tree for each source and its destination host group. The distribution tree provides a shortest path between the source and each multicast receiver in the group, based on the number of hops in the path. A tree is constructed on demand, using a broadcast and prune technique, when a source begins to transmit messages to a multicast group.

DVMRP assumes that initially every host on the network is part of the multicast group. The designated router on the source subnet (the router that is selected to handle routing for all hosts on the subnet) begins transmitting a multicast message to all adjacent routers. Each of these routers selectively forwards the message to downstream routers until the message is eventually passed to all multicast group members.

Navigation

- Reverse path forwarding on page 42
- Pruning and grafting on page 42
- DVMRP concepts and terminology on page 43
- DVMRP static source groups on page 45
- DVMRP routing policies on page 45
- Avaya Ethernet Routing Switch 8800/8600 implementation of DVMRP on page 51

Reverse path forwarding

In the selective forwarding process during the formation of the multicast tree, after a router receives a multicast stream, it checks the DVMRP routing tables to determine the interface that provides the shortest path back to the source. If the shortest path is from the interface where the multicast stream arrived, the router enters state information to identify the multicast stream and the source in the internal tables, and forwards the multicast message to all adjacent routers except those on the same interface. If the interface that receives the multicast stream is not the shortest path, the stream is discarded. This mechanism, called reverse path forwarding, ensures that the tree does not contain loops and that the tree includes the shortest path from the source to all recipients.

Pruning and grafting

Pruning eliminates branches of the distribution tree that do not lead to multicast group members. The IGMP running between hosts and their immediately neighboring multicast routers is used to maintain group membership data in the routers. After a router determines that no hosts beyond it belong to the multicast group, it sends a prune message to its upstream

router. Routers update source and destination group state information in their tables to reflect the branches that are eliminated from the tree, resulting in a minimum multicast tree. If a router later learns of new group memberships from the hosts or downstream routers, it sends a graft message upstream to retract the prune sent earlier.

After the multicast tree is constructed, it is used to transmit multicast messages from the source to multicast members. Each router in the path forwards messages over only those interfaces that lead to group members. Because new members can join the group and these members can depend on one of the pruned branches to receive the transmission, DVMRP periodically reinitiates the construction of the multicast tree.

DVMRP concepts and terminology

DVMRP provides a mechanism for routers to propagate multicast datagrams in a manner that minimizes the number of excess copies sent to a particular network.

Neighbor connections

In a DVMRP environment, neighbors are multicasting routers that use an interface to the same network.

At startup, a DVMRP multicasting router performs the following tasks:

- initializes its routing table with information about all of its local networks
- learns the existence of its neighbors by sending a probe for all routes on each of its multicast interfaces
- receives reports from its neighbors containing the routing information (including route costs)

Source route advertisements

A source network is a network containing hosts that can issue multicast datagrams. DVMRP advertises shortest-path routes to multicasting source networks.

Periodically, each multicasting router issues full or partial routing information on each DVMRP circuit using DVMRP report messages. This routing information represents the cost for the sending router to reach the specified source network. The cost is the sum of the hop metrics along the shortest path to the source network.

Upon receiving a DVMRP report from another router, DVMRP reexamines its routing table to determine whether the shortest path information needs updating. Specifically, DVMRP looks in the routing table for an entry describing a route to the same source network. If one exists, DVMRP compares the cost of the two routes and stores the route with the lower cost in its routing table.

A router does not send route reports on an interface until it knows (by means of received probes or reports) that a neighboring multicast router exists on that interface. The Avaya Ethernet Routing Switch 8800/8600 acknowledges implicit probes from neighboring multicast routers, and it sends probes periodically on the interface.

How DVMRP chooses a route

Each DVMRP interface is configured with a metric that indicates the cost of the hop. A router that receives multiple route reports for the same multicasting source network performs the following tasks:

- compares the cost specified in each route report (based on the metric field)
- stores information from the report with the lowest cost in its routing table

A route metric is the sum of all the interface (hop) metrics from a route source to a router. After a next-hop neighbor is declared for a route, the route updates received from that neighbor for that route take precedence until either the route times out or another router advertises a better metric for that route.

Routing table

Table 2: Parts of a routing table entry on page 44 shows the principal items in a routing table entry.

Table 2: Parts of a routing table entry

Item	Description
Source subnet address and mask	The network address and mask that identify that the source for this entry contains multicast routing information.
Upstream neighbor	The address of the upstream neighbor from where multicast datagrams are received.
Interface	The value of the interface index where IP datagrams sent by these sources are received.
Metric	The distance in hops to the source subnet.
Expiration Time	The maximum amount of time (in time ticks) remaining before this entry ages out.

The source subnet and the previous-hop router in the DVMRP routing table are the opposite of the destination subnet and next-hop router in a RIP routing table.

The router uses this information to perform the following tasks:

- receive a multicast datagram and determines whether it arrived on the interface that is on the shortest path to the source network
- drop the datagram if it did not arrive on the shortest-path interface
- flood the multicast stream to all active, nonpruned, downstream DVMRP neighbors

Shortest-path trees

Route information used by DVMRP is independent of other routing information used by the router. DVMRP routing information creates a shortest-path tree entry in the routing table for the propagation of multicast datagrams.

The shortest-path tree entry indicates the interface that provides the shortest path to the network that is the source of the multicast datagram. A shortest-path tree also indicates those interfaces that are on the shortest path to that source network from a neighboring router.

In IGMPv2, neighboring routers use the same metric to a source network. The router with the lower IP address propagates multicast traffic originating from that source network onto the network or tunnel that is common to these neighboring routers.

A network is considered a leaf network if it does not depend on downstream neighbors for a source.

DVMRP static source groups

Static source groups enable you to configure static source group entries in the DVMRP multicast routing table. DVMRP cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 25.

DVMRP routing policies

With DVMRP routing policies, you can improve the management of the DVMRP routing tables by controlling how the routing table is populated and how the routes are exchanged between switches. You can apply these routing policies, when enabled, to either a VLAN or a brouter port.

Default route policy

DVMRP uses a default route to summarize routes in the routing table to reduce the size of the routing table, which is particularly useful for the edge switches in your network. DVMPR uses default route policies DVMRP to perform the following tasks:

- Listen for a default route: You can enable or disable a switch to listen for a default route; the Avaya Ethernet Routing Switch 8800/8600 is configured by default to listen for a default route.
- Advertise a default route: You can enable or disable a switch to advertise a default route; the Avaya Ethernet Routing Switch 8800/8600 is configured by default to advertise a default route if the route exists in the routing table.
- Supply a default route: You can configure an interface to supply a default route, where the default route is generated and advertised. In this case, the default route is not added to the routing table but is used by a neighbor switch as a path from which all unknown source addresses are accepted.

When you supply a default route, consider the following:

- You can set the Avaya Ethernet Routing Switch 8800/8600 to advertise a default route only if that switch is not enabled to listen for the default route.
- If you configure an interface to supply a default route, it does not advertise other routes to its neighbors.
- The Avaya Ethernet Routing Switch 8800/8600 does not listen for default routes on an interface that is configured to supply a default route.

Figure 6: Default route configuration example on page 46 shows a network with two domains that include multiple switches. In this configuration, S1 is an edge switch that connects only to S4 and S5 in a different domain. S4 and S5 are configured to supply default routes to S1. The metrics for these switches are set as one hop for S1 and two hops for S5. S4 and S5 do not advertise other routes to S1 because they are configured to supply the default route. In this case, the default route is not added to the S4 and S5 routing tables because the interface supplies default routes.

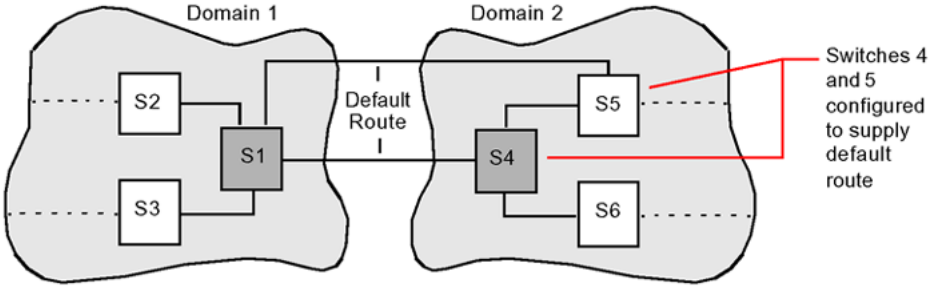

Figure 6: Default route configuration example

In this example, only one default route is active at a time; therefore, no load sharing occurs on the links. In addition, in some cases, a nonoptimal path is taken to reach S1 from a switch in

the other domain. By default, the other switches in the domain that includes S1 receive the default route advertised by S1 the same as a regular route. No exceptions exist as none of the switches are configured with a supply default route policy. If required, configure these switches to not accept the route.

All the switches with a default route in their routing table accept traffic on the interface where they learn the default route from a source that is unknown to them.

You can simplify the configuration shown in Figure 6: Default route configuration example on page 46. In the simplified configuration, S1 uses only one link to domain 2, and S4 is configured to supply a default route to S1. In this configuration, S1 accepts all multicast traffic from domain 2 on that interface.

For more information about configuration examples that show the application of the DVMRP announce and accept policies to the same network configuration shown in Figure 6: Default route configuration example on page 46, see Figure 8: Announce policy configuration example on page 48 and Figure 10: Accept policy configuration example on page 50.

Announce policy

A DVMRP announce policy (out filter) governs the propagation of DVMRP routing information. Use DVMRP announce policies to control which routes are sent to neighboring routers, to reduce the size of routing tables, and provide a level of security for the network.

The announce policies apply to the outgoing advertisements to the other neighbors or peers in the protocol domain. These policies determine whether to announce specific route information. You can selectively announce routes with announce policies. You can configure a policy to apply to a route. If no configured policy or matching policy for a route exists, the default configuration accepts and announces the route.

Figure 7: DVMRP announce policy logic on page 48 shows what happens to an outgoing route when a DVMRP announce policy does or does not exist on a switch.

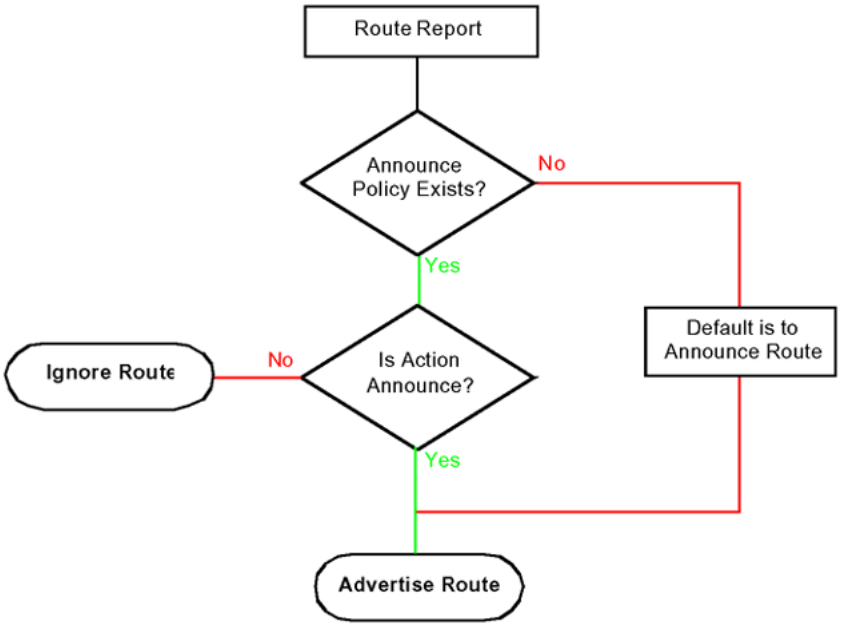

Figure 7: DVMRP announce policy logic

Figure 8: Announce policy configuration example on page 48 shows an announce policy communication between two different domains that include multiple switches. In this example, S1 is an edge switch that connects only to S4 and S5, which are in a different domain. S4 is configured so that it does not announce routing information. The result and benefit of this configuration is that the local switch (S4) does not send route advertisements to the other switches in the network, therefore, reducing the size of the routing tables and providing a level of security for the network.

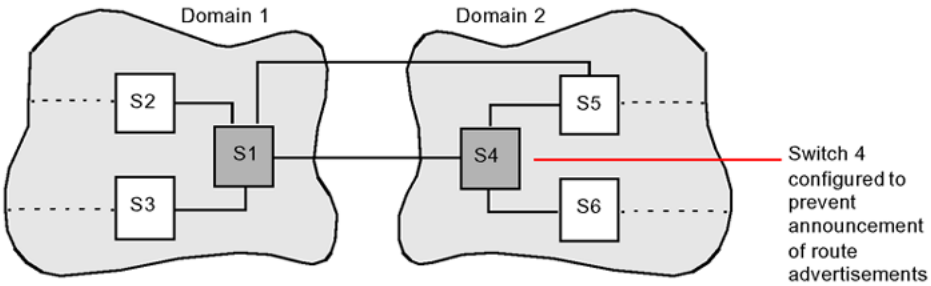

Figure 8: Announce policy configuration example

Accept policy

A DVMRP accept policy (in filter) controls the way DVMRP manages incoming route advertisements. Accept policies apply to incoming advertisements and reduce the size of the DVMRP routing table. For example, you can configure an edge switch to use a default route and not to accept other routes, therefore, reducing the size of its routing table (the routing table

includes only the default and local routes). In this case, this switch can still advertise all of its routes to the rest of the network.

Accept policies inject routes into the DVMRP routing table; you can apply them to single or all interfaces of a switch. Configure a policy to apply to a specific router to selectively accept routes. If no configured policy or matching policy for a route exists, the default configuration accepts the route.

Figure 9: DVMRP accept policy logic on page 49 shows what happens to an incoming route when a DVMRP accept policy is or is not enabled on your switch.

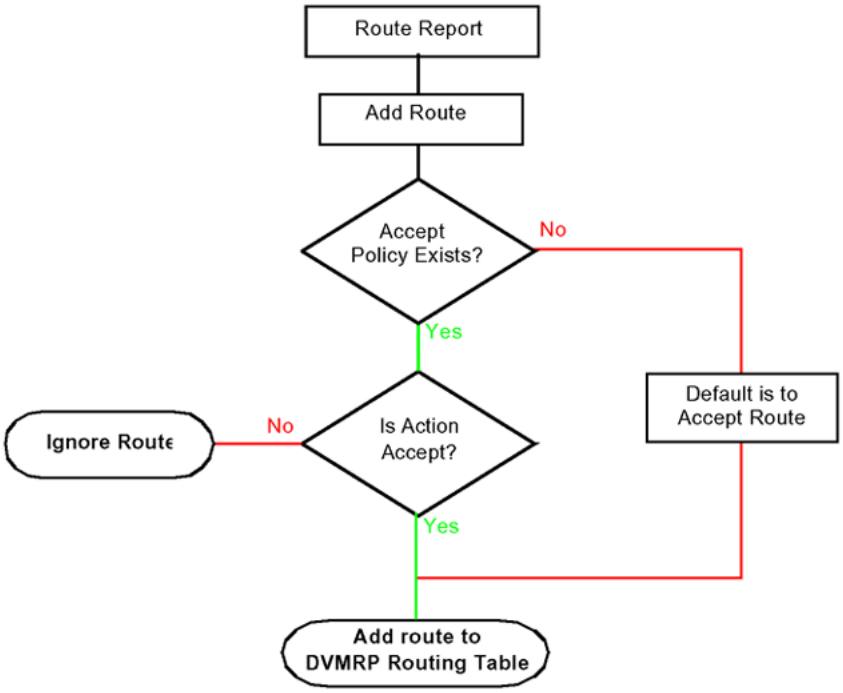

Figure 9: DVMRP accept policy logic

Figure 10: Accept policy configuration example on page 50 shows an accept policy communication between two different domains that include multiple switches. In this example, S1 is an edge switch that connects only to S4 and S5, which are in a different domain. S4 and S5 are configured so that they do not accept routing information from S1. The result and benefit of this configuration is that the switches in domain 2 do not receive routing information from domain 1, therefore, reducing the size of the routing tables and providing a level of security for the network. S4 and S5 do, however, advertise routing information to S1. Therefore, switches in domain 1 receive the routing information from domain 2.

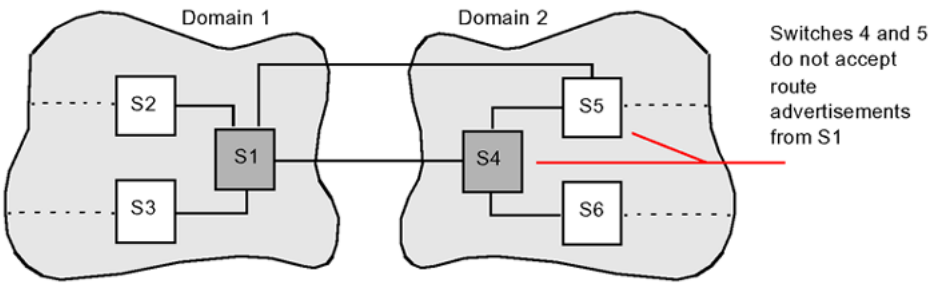

Figure 10: Accept policy configuration example

Advertisement of local networks policy

By default, DVMRP advertises its own local networks over an interface. With the advertisement of local networks policy, you can configure an interface to omit the advertisement of local routes to other switches in your network. This policy reduces the size of the routing table and provides a level of security where multicast traffic on the interface is not routed to other interfaces in the network. The interface still receives multicast traffic from the other interfaces in the network.

Figure 11: Advertisement of local networks policy configuration example on page 50 shows how the advertisement of a local networks policy applies to switches communicating with a core network. In this example, four switches with several interfaces contain receivers (such as a television) that communicate with devices in the core network. The switches are configured to disable the advertisement of their local network, which means that the switch only receives (does not supply) multicast traffic. With this configuration, the switch maintains smaller routing tables.

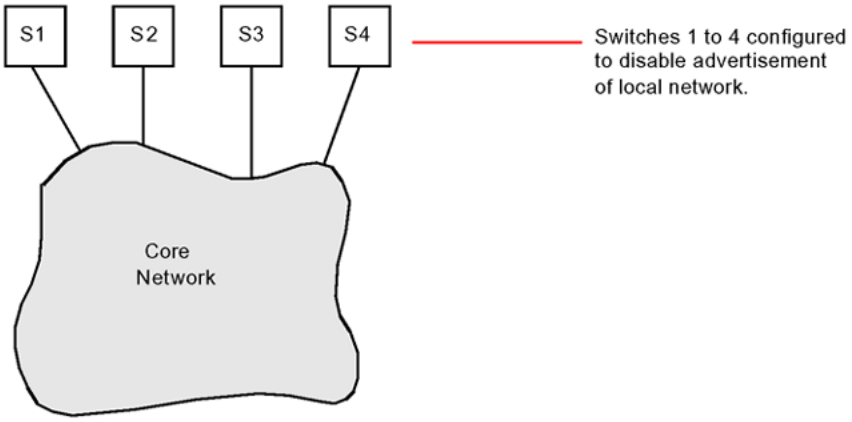

Figure 11: Advertisement of local networks policy configuration example

DVMRP passive interface policy

With the DVMRP passive interface policy, you can configure multiple DVMRP interfaces on an Avaya Ethernet Routing Switch 8800/8600 without affecting the performance of the switch. You can configure an interface as passive or active. If you configure an interface as passive,

it drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to its neighbor switches. If you configure a DVMRP interface as passive, you can only change the interface type if the interface is disabled.

Avaya Ethernet Routing Switch 8800/8600 implementation of DVMRP

In an Avaya Ethernet Routing Switch 8800/8600, DVMRP fully supports multiaccess networks. The forwarding entries for the receivers on multiaccess networks are port based rather than network based. Therefore, on a multiaccess network, only ports interested in the data receive it. That is, IP multicast routing is supported on ports with port-based or IP subnet-based VLANs enabled.

The DVMRP router listens to all IGMP host membership reports even if it is not the designated querier and keeps a local group database of every host membership reporter.

When a multicast stream of UDP packets first enters the switch, if DVMRP is enabled for the interface, it processes these packets as necessary and creates a hardware cache entry to handle subsequent packets in the same stream for the same multicast destination. The packets are discarded if no members exist; otherwise they are forwarded.

The Avaya Ethernet Routing Switch 8800/8600 implementation does not support DVMRP tunneling.

Protocol Independent Multicast-Sparse Mode

PIM-SM, as defined in RFC 2362, supports multicast groups spread out across large areas of a company or the Internet. Unlike dense-mode protocols, such as DVMRP, which initially flood multicast traffic to all routers over an entire internetwork, PIM-SM sends multicast traffic only to routers that specifically joined a multicast group. This technique reduces traffic flow over WAN links and overhead costs for processing unwanted multicast packets.

Dense-mode protocols use a flood-and-prune technique, which is efficient where receivers are densely populated. However, for sparsely populated networks, PIM-SM is more efficient because it sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic.

PIM-SM is independent of a specific unicast routing protocol, but it does require the presence of a unicast routing protocol, such as RIP, Open Shortest Path First (OSPF), or static routing. PIM-SM uses the information from the unicast routing table to create and maintain multicast trees that enable PIM-enabled routers to communicate.

Navigation

- PIM-SM concepts and terminology on page 52
- Shared trees and shortest-path trees on page 57
- Receiver joining a group on page 58
- Receiver joining a group on page 58
- Source sending packets to a group on page 59
- Required elements for PIM-SM operation on page 60
- PIM-SM simplified example on page 60
- PIM-SM static source groups on page 61
- PIM-SMLT on page 62

PIM-SM concepts and terminology

Typically, a PIM-SM network consists of several multipoint data streams, each targeted to a small number of LANs in the internetwork. For example, customers whose networks consist of multiple hosts on different LANs can use PIM-SM to simultaneously access a video data stream, such as video conferencing, on a different subnet.

Important:

In some cases, PIM stream initialization can take several seconds.

Hosts

A host is a source, a receiver, or both:

- A source, also known as a sender, sends multicast data to a multicast group.
- A receiver receives multicast data from one or several sources that sends data to a multicast group.

PIM-SM domain

PIM-SM operates in a domain of contiguous routers on which PIM-SM is enabled. These routers are configured to operate within a common boundary defined by a PIM multicast border router (MBR).

Each PIM-SM domain requires the following routers:

- designated router (DR)
- rendezvous point router
- bootstrap router (BSR)

Although a PIM-SM domain can use only one active RP router and one active BSR, you can configure additional routers as a candidate RP (C-RP) router and as a candidate BSR (C-BSR). Candidate routers provide backup protection in case the primary RP router or BSR fails.

Designated router

The DR is the router with the highest IP address on a LAN designated to perform the following tasks:

- sends register messages to the RP router on behalf of directly connected sources
- sends join and prune messages to the RP router on behalf of directly connected receivers
- maintains information about the status of the active RP router for local sources in each multicast group

Important:

The DR is not a required configuration, and switches act automatically as the DR for directly attached sources and receivers.

Rendezvous point router

PIM-SM builds a shared multicast distribution tree within each domain, and the RP router is at the root of this shared tree. Although the RP can physically exist anywhere on the network, it must be as close to the source as possible. Only one active RP router exists for a multicast group.

At the RP router, receivers meet new sources. Sources use the RP to identify themselves to other routers on the network; receivers use the RP to learn about new sources.

The RP performs the following tasks:

- registers a source that wants to announce itself and send data to group members
- joins a receiver that wants to receive data for the group
- forwards data to group

Candidate rendezvous point router

You can configure a set of routers as C-RP routers that serve as backup to the RP router. If an RP fails, all the routers in the domain apply the same algorithm to elect a new RP from the group of C-RP routers. To make sure that the routers use a complete list of C-RP routers, the C-RP router periodically sends unicast advertisement messages to the BSR. The most

common implementation is to configure a PIM-SM router as both a C-RP router and a C-BSR.

Important:

Although you can configure a C-RP router on a DVMRP interface, no functionality is tied to this configuration.

Static rendezvous point router

You can configure a static entry for a rendezvous point (RP) with static RP. This feature avoids the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. Static RP-enabled switches cannot learn about RPs through the BSR because the switch loses all dynamically learned BSR information and ignores BSR messages. After you configure static RP entries, the switch adds them to the RP set as if they were learned through the BSR.

Important:

In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interface configured as an RP.

When you configure a PIM static RP in a switch, the next hop of the unicast route toward PIM static RP must be a PIM neighbor. The PIM protocol fails to work, if due to a route change, the next hop toward an already configured static RP becomes a non-PIM neighbor. The configured RP does not activate until it is reachable through a PIM neighbor, and its state remains invalid.

A static RP-enabled Avaya Ethernet Routing Switch 8800/8600 can communicate with switches from other vendors that do not use the BSR mechanism. Some vendors use either early implementations of PIM-SM v1 that do not support the BSR or proprietary mechanisms like the Cisco Auto-RP. For a network to work properly with static RP, you must map all the switches in the network (including switches from other vendors) to the same RP or RPs, if several RPs are present in the network.

To avoid a single point of failure, you can also configure redundant static RPs.

Use the static RP feature when you do not need dynamic learning mode, typically in small networks, or for security reasons, where you force RPs to some devices in the network so that they do not learn other RPs.

Specific route for static RP

With static RP enabled, the Avaya Ethernet Routing Switch 8800/8600 detects RP failure based on IGP convergence and, more specifically, on the removal of the route to the RP from the routing table. With the route to the failed RP removed, the Avaya Ethernet Routing Switch 8800/8600 can fail over to an alternate static RP.

If a default route is injected into the routing table, that default route still appears as an active route to the failed RP. Therefore, in this case, the switch does not fail over to the alternate RP.

A similar situation exists with SMLT-based configurations, where an internal-only default static route is used during IST failover and recovery. In this case, the internal default route appears as an active route to the failed RP, and therefore does not failover to the alternate RP.

To resolve the preceding situations, you can configure the lookup for static RP to be chosen from the specific route rather than the best route. In this case, when the route to the active RP fails, the switch no longer interprets the default route as a valid route for RP purposes, and therefore fails over to the alternate RP.

Avaya recommends that you always enable the specific route option for any SMLT/RSMLT cluster running PIM-SM with static RPs because of the implementation of the internal-only default static route on the IST.

This resolution applies only to static RP configurations, not to C-RP configurations.

RP Set configuration considerations

When you configure RP sets (C-RPs or static RPs), Avaya recommends as best practice not to configure multiple entries that each specify a unique group, but instead specify a range of groups when possible, thereby decreasing the number of entries required.

Static RP configuration considerations

Before you can configure a static RP, you must enable PIM in Sparse Mode (SM) and enable static RP.

After you meet these prerequisites, keep in mind the following configuration considerations:

- You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.
- All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.
- Static RPs do not age, that is, they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless it is configured with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP.
- In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interfaces configured as an RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of Avayal and other vendor switches across the network, you must ensure that all switches and routers use the same active RP because other vendors can use different algorithms to elect the active RP. Avaya Ethernet Routing Switch 8800/8600 devices use the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.

Important:

To reduce convergence times, Avaya recommends that you create only one static RP for each group. The more static RPs you configure for redundancy, the more time PIM requires to rebuild the mroute table and associate RPs.

- Static RP configured on the switch is active as long as the switch uses a unicast route to the static RP network. If the switch loses this route, the static RP is invalidated and the hash algorithm remaps all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm remaps the affected groups.

Bootstrap router

The BSR receives RP router advertisement messages from the candidate RPs. The BSR adds the RP router with its group prefix to the RP set. Only one BSR exists for each PIM-SM domain.

The BSR periodically sends bootstrap messages containing the complete RP set to all routers in the domain. The BSR ensures that all PIM-SM routers send join/prune and register packets.

Candidate bootstrap router

Within a PIM-SM domain, you can configure a small set of routers as C-BSRs. The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Important:

Configure C-BSRs on routers that are central to all candidate RPs.

Join and prune messages

The DR sends join and prune messages from a receiver toward an RP for the group to either join the shared tree or remove (prune) a branch from it. A single message contains both a join and a prune list. This list includes a set of source addresses indicating the shortest-path trees or the shared trees that the host wants to join. The DR sends join and prune messages hop by hop to each PIM router on the path to the source or the RP.

Register and register-stop messages

The DR sends register messages to the RP for a directly connected source. The register message informs the RP of a new source, causing the RP to send join or prune messages back toward the DR of the source, which forwards the data down the RP tree after it gets the

data natively. After the receiver DR gets the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.

The DR stops sending encapsulated packets to the RP after receiving a register-stop message. This traffic stops without delay because the RP sends a register-stop message immediately after receiving the first multicast data packet, and joins the shortest-path tree.

Shared trees and shortest-path trees

In a PIM-SM domain, shared trees and shortest-path trees are used to deliver data packets to group members. This section describes both trees.

Shared trees

Group members in a PIM-SM domain receive the first packet of data from sources across a shared tree. A shared tree consists of a set of paths that connect all members of a multicast group to the RP. PIM creates a shared tree when sources and receivers send messages toward the RP.

Shortest-path trees

After receiving a certain number of packets from the RP, the DR switches from a shared tree to a shortest-path tree (SPT). Switching to an SPT creates a direct route between the receiver and the source. The Avaya Ethernet Routing Switch 8800/8600 switches to the SPT after it receives the first packet from the RP.

Figure 12: Shared tree and shortest-path tree on page 58 shows a shared tree and an SPT.

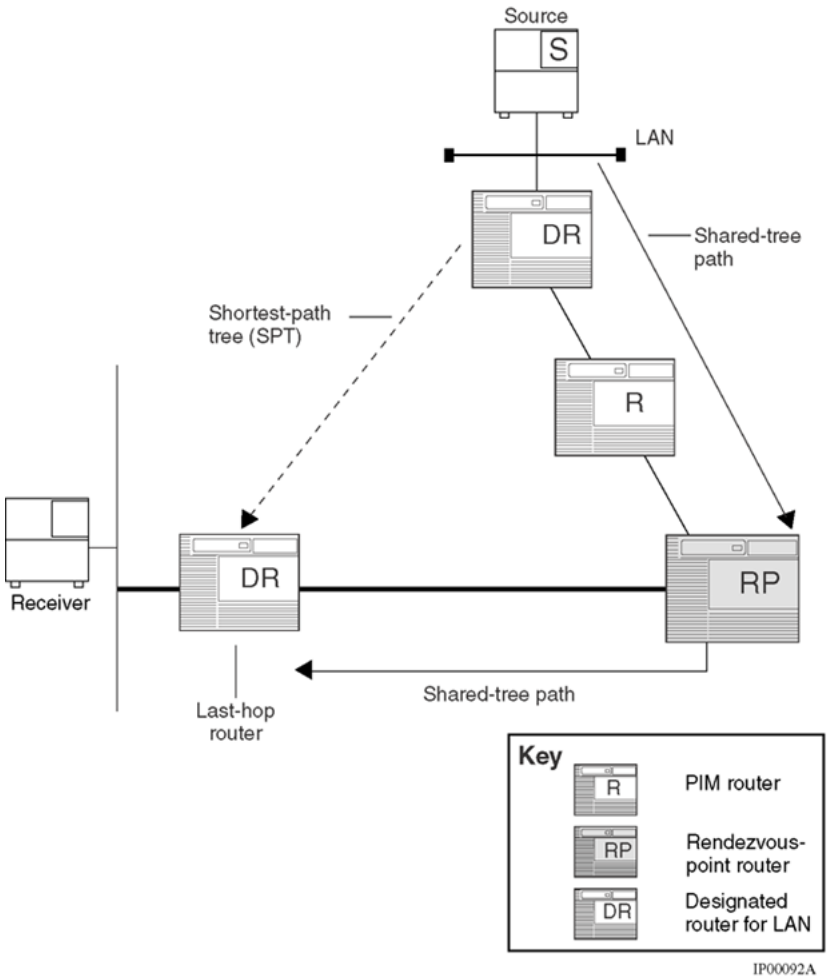

Figure 12: Shared tree and shortest-path tree

Receiver joining a group

The following steps describe how a receiver joins a multicast group:

1. A receiver multicasts an IGMP host membership message to the group that it wants to join.
2. After the last-hop router (the DR), which is normally the PIM router with the highest IP address for that VLAN, receives the IGMP message for a new group join, the router looks up the associated elected RP with responsibility for the group.
3. After it determines the RP router for the group, the last-hop router creates a (*,G) route entry in the multicast forwarding table and sends a (*,G) join message to the RP. After the last-hop router receives data packets from the RP, if the multicast packet arrival rate exceeds the DR threshold, the last-hop router switches to the

SPT by sending an (S,G) join message to the source. (S denotes the source unicast IP address, and G denotes the multicast group address.)

4. If the switch to the SPT occurs
   - All intermediate PIM routers along the path to the source create the (S,G) entry.
   - To trim the shared tree, the router sends an (S,G) prune message to the RP.

Receiver leaving a group

Before it leaves a multicast group, a receiver sends an IGMP leave message to the DR. If all directly connected members of a multicast group leave or time out, and no downstream members remain, the DR sends a prune message upstream and PIM-SM deletes the route entry after that entry times out.

Source sending packets to a group

The following steps describe how a source sends multicast packets to a group:

1. A source directly attached to a VLAN bridges the multicast data to the DR. The DR for the VLAN (the router with the highest IP address) encapsulates each packet in a register message and sends a unicast message directly to the RP router to distribute to the multicast group.
2. If a downstream group member chooses to receive multicast traffic, the RP router sends a join/prune message toward the source DR and forwards the data down the RP tree after it gets the data natively.
3. After the receiver DR gets the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.
4. If no downstream members want to receive multicast traffic, the RP router sends a register-stop message (for the source) to the DR.

   The DR starts the register suppression timer after it receives the first register-stop message. During the register suppression timeout period (the default is 60 seconds), the following events occur:

   - The DR for the source sends a probe packet to the RP router before the register suppression timer expires. The probe packet prompts the RP router to determine whether new downstream receivers joined the group.
   - If no new receivers joined the group, the RP router sends another register-stop message to the DR for the source, and its register suppression timer restarts.
   - If the RP router no longer responds with a register-stop message to the source DR probe message, the register suppression timer expires and the DR sends

encapsulated multicast packets to the RP router. The RP router uses this method to tell the DR that new members joined the group.

The RP sends a register-stop message to the DR immediately after receiving the first multicast data packet.

Required elements for PIM-SM operation

For PIM-SM to operate, the following elements must exist in the PIM-SM domain:

- You must enable an underlying unicast routing protocol for the switch to provide routing table information to PIM-SM.
- You must configure an active BSR to send bootstrap messages to all PIM-v2 configured switches and routers to enable them to learn group-to-RP mapping. If several BSRs are configured in a network, an active BSR is elected based on priority and IP address (if priority is equal, the BSR with the higher IP address is elected).
- You must include an RP to perform the following tasks:
  - manage one or several IP multicast groups
  - become the root for the shared tree to these groups
  - accept join messages from receiver switches for groups that it manages
  - elect an active RP based on priority and IP address (if priority is equal, the RP with the higher IP address is elected)

PIM-SM simplified example

Figure 13: PIM-SM simplified example on page 61 shows a simplified example of a PIM-SM configuration.

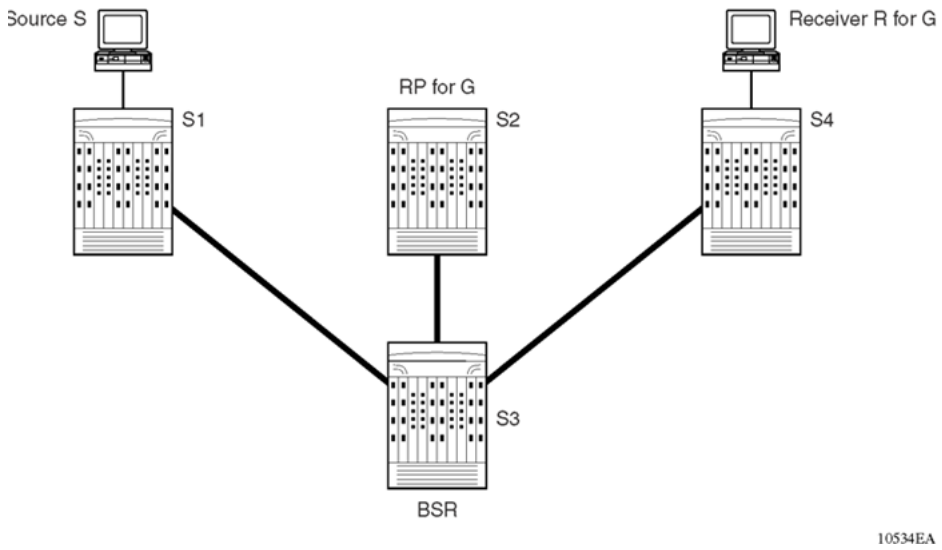

Figure 13: PIM-SM simplified example

In the sample configuration, the following events occur:

1. The BSR distributes RP information to all switches in the network.
2. R sends an IGMP membership report to S4.
3. Acting on this report, S4 sends a (*,G) join message to RP.
4. S sends data to G.
5. The DR (S1 in this example) encapsulates the data that it unicasts to RP (S2) in register messages.
6. S2 decapsulates the data, which it forwards to S4.
7. S4 forwards the data to R.
8. If the packet rate exceeds the DR threshold, S4 sends S1 an (S,G) join message.
9. S1 forwards data to S4. After S4 receives data from S1, it prunes the stream from the RP.

Important:

Figure 13: PIM-SM simplified example on page 61 is a simplified example and is not the best design for a network if the source and receiver are placed as shown. In general, RPs are placed as close as possible to sources.

PIM-SM static source groups

You can configure static source groups as static source-group entries in the PIM-SM multicast routing table. PIM-SM cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 25.

PIM-SMLT

IP multicast routing support with SMLT builds a virtual switch that represents the two switches of the split multilink trunk core. When switches run PIM in the core, they need to exchange protocol-related updates as part of the interswitch trunking (IST) protocol. IST hides the fact that the edge switch is attached to two physical switches.

Important:

DVMRP is not supported in the core as a Layer 3 routing protocol.

PIM-SMLT works in triangular, square, and full-meshed configurations with Layer 3 IP multicast.

The following rules apply:

- For a VLAN, if traffic is received from the IST link, it cannot forward on the split multilink trunk link or the edge for the same VLAN.
- Traffic can use the IST to route between VLANs if the forwarding decision for the multicast protocol requires that the other side of the core forwards the multicast traffic (follow the IP multicast routing and forwarding rules for routed traffic). Other VLANs that are not part of SMLT continue to behave in the same way.
- To create a temporary default route pointing to a peer IST, you must enable PIM on the IST VLAN.

SMLT provides for fast failover in all cases, but is not designed to provide a functionality similar to RSMLT.

Important:

You must enable square SMLT globally before configuring square or full-mesh configurations.

Traffic delay with PIM while rebooting peer SMLT switches

PIM uses a DR to forward data to receivers on the VLAN. If you reboot the DR in an SMLT VLAN, you can lose data because of the following actions:

- When the DR is down, the non-DR switch assumes the role and starts forwarding data.
- After the DR comes back up, it takes priority (higher IP address) to forward data so the non-DR switch stops forwarding data.
- The DR is not ready to forward traffic due to protocol convergence and because it takes time to learn the RP set and create the forwarding path. This can result in a traffic delay of 2 to 3 minutes because the DR learns the RP set after Open Shortest Path First (OSPF) converges.

A workaround is to a configure static RP router on the peer SMLT switches. This feature avoids the process of selecting an active RP router from the list of candidate RPs and dynamically

learning about RPs through the BSR mechanism. After the DR comes back up, traffic resumes as soon as OSPF converges. This workaround reduces the traffic delay to approximately 15 to 65 seconds.

PIM-SSM

Source Specific Multicast optimizes PIM-SM by simplifying the many-to-many model. Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM only builds source-based SPTs. Whereas PIM-SM always joins a shared tree first, and then switches to the source tree, SSM eliminates the need to start with a shared tree by immediately joining a source through the SPT. SSM avoids using a rendezvous point (RP) and RP-based shared trees, which can be a potential problem.

Members of an SSM group can only receive from a single source. This configuration is ideal for applications like television channel distribution and other content-distribution businesses. Banking and trade applications can also use SSM because it provides more control over the hosts receiving and sending data over their networks.

SSM applications use IP addresses reserved by the Internet Assigned Numbers Authority (IANA) in the 232/8 range (232.0.0.0 to 232.255.255.255). SSM recognizes packets in this range and controls the behavior of multicast routing devices and hosts that use these addresses.When a source (S) transmits IP datagrams to an SSM destination address (G), a receiver can subscribe to the (S,G) channel to receive these datagrams.

A channel is a source-group (S,G) pair where S is the source sending to the multicast group and G is an SSM group address. SSM defines channels on an individual source basis, which enforces the one-to-many concept of SSM applications. In an SSM channel, each group is associated with only one source. However, another SSM channel can associate the same multicast group with a different source, which provides an efficient use of the SSM address range. For example, channel 192.1.3.4, 232.1.2.3 is different from channel 141.251.186.13, 232.1.2.3.

SSM features

SSM uses only a subset of the PIM-SM features such as the SPT, DR, and some messages (hello, join/prune, and assert). However, some features are unique to SSM. These features, which are described in the following sections, are extensions of the IGMP and PIM protocols.

PIM requires a unicast protocol to forward multicast traffic within the network when it performs the Reverse Path Forwarding (RFP) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that PIM-enabled

routers use to communicate. The unicast routing table must contain a route to every multicast source in the network as well as routes to PIM entities like the RPs and BSR.

PIM-SSM architecture

The following figure illustrates how the PIM-SSM architecture requires routers to

- support IGMPv3 source-specific host membership reports and queries at the edge routers
- initiate PIM-SSM (S,G) joins directly and immediately after receiving an IGMPv3 join report from the designated router
- restrict forwarding to SPTs within the SSM address range by all PIM-SSM routers

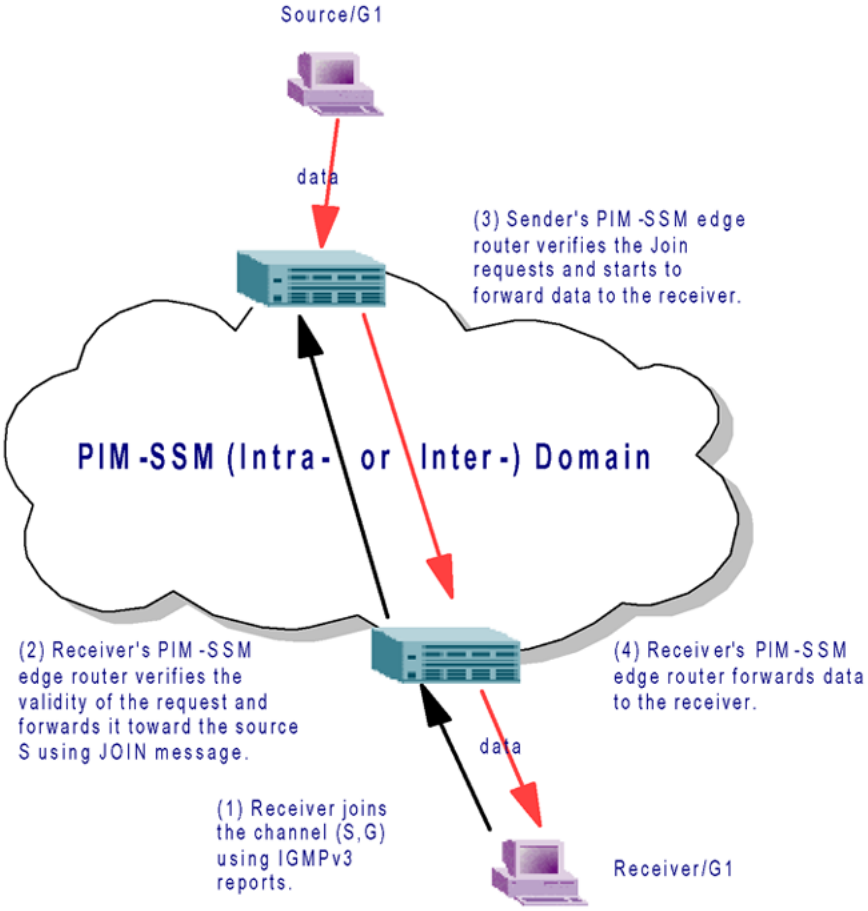

Figure 14: PIM-SSM architecture

The following rules apply to Layer 3 devices with SSM enabled:

- Receive IGMPv3 membership join reports in the SSM range and, if no entry (S,G) exists in the SSM channel table, create one.
- Receive IGMPv2 membership join reports, but only for groups that already use a static (S,G) entry in the SSM channel table.
- Send periodic join messages to maintain a steady SSM tree state.
- Use standard PIM-SM SPT procedures for unicast routing changes, but ignore rules associated with the SPT for the (S,G) route entry.
- Receive prune messages and use standard PIM-SM procedures to remove interfaces from the source tree.
- Forward data packets to interfaces from the downstream neighbors that sent an SSM join, or to interfaces with locally attached SSM group members.
- Drop data packets that do not use an exact-match lookup (S,G) in their forwarding database for S and G.

Avaya Ethernet Routing Switch 8800/8600 implementation of SSM and IGMP

The following sections describe how PIM-SSM and IGMP are implemented in the Avaya Ethernet Routing Switch 8800/8600.

SSM range

The standard SSM range is 232/8, but you can extend the range to include an IP multicast address. Although you can configure the SSM range, you cannot configure it for all multicast groups (224/4 or 224.0.0.0/240.0.0.0 or 224.0.0.0/255.0.0.0).

You can extend the SSM range to configure existing applications without changing their group configurations.

SSM channel table

You can use the SSM channel to manually configure S,G entries that map existing groups to their sending source. These table entries apply to the whole switch, not for each interface, and both IGMPv2 and IGMPv3 hosts use the SSM channel table.

The following rules apply to an SSM channel table for an individual switch:

- You can map one source to multiple groups.
- You can map one group to one source only; that is, you cannot map the same group more than once in a table.

Important:

Different switches can use different mappings for groups to sources, for example, different channels are mapped differently even if they are on the same network.

SSM and IGMPv2

SSM-configured switches can accept reports from IGMPv2 hosts on IGMPv2 interfaces if the group uses an SSM channel table entry. However, the IGMPv2 host groups must exist in the SSM range defined on the switch, which is 232/8 by default.

- After the SSM switch receives an IGMPv2 report for a group that is in the SSM channel table, it joins the specified source immediately.
- After the SSM switch receives an IGMPv2 report for a group that uses an enabled static SSM channel table entry, it triggers PIM-SSM processing as if it received an equivalent IGMPv3 report.
- After the SSM switch receives an IGMPv2 report for a group out of the SSM range, it processes the report as if it is in PIM-SM mode.

SSM and IGMPv3

The Avaya Ethernet Routing Switch 8800/8600 supports IGMPv3 for SSM. With IGMPv3, a host can selectively request or filter traffic from sources within the multicast group. IGMPv3 is an interface-level configuration.

Important:

IGMPv3 works only with PIM-SSM or SSM-snoop enabled on the interface.

The following rules apply to IGMPv3-enabled interfaces:

- Send only IGMPv3 (source-specific) reports for addresses in the SSM range.
- Accept IGMPv3 reports.
- Drop IGMPv2 reports.

  The IGMPv2 report mentioned in SSM and IGMPv2 on page 66 is processed because it is an IGMPv2 report that is received on an IGMPv2 interface. If an IGMPv2 interface

receives an IGMPv3 report, the report is dropped even if PIM-SSM is enabled and the entry is in the SSM channel table. The IGMP version must match.

- Discard IGMP packets with a group address out of the SSM range.

The Avaya Ethernet Routing Switch 8800/8600 implements IGMPv3 in one of two modes: dynamic and static.

- In dynamic mode, the switch learns about new (S,G) pairs from IGMPv3 reports and adds them to the SSM channel table. If dynamic mode is not enabled and an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report.
- In static mode, you can statically configure (S,G) entries in the SSM channel table. If an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report. It also ignores the report if the group is in the table, but the source or mask does not match what is in the table.

Important:

When you enable IGMPv3, changes to the query interval and robustness values on the querier switch propagate to other switches on the same VLAN through IGMP query.

Both IGMPv2 and IGMPv3 hosts use the SSM channel table:

- An IGMPv2 host (with an IGMPv2 VLAN) must use an existing SSM channel entry if the group is in the SSM range.
- If you enable dynamic learning for an IGMPv3 host, the SSM channel automatically learns the group. Otherwise, the SSM channel also needs a static entry.

The following table summarizes how a switch in PIM-SSM mode works with IGMP. References to matching a static SSM channel entry assumes that the entry is enabled. If an entry is disabled, it is treated as though it is disallowed.

Table 3: Summary of PIM-SSM interaction with IGMPv2 and v3

Host	VLAN	SSM range	Action
IGMPv2 host	IGMPv3 VLAN	In or out of range	Drop report.
IGMPv3 host	IGMPv2 VLAN	In or out of range	Drop report.
IGMPv2 host	IGMPv2 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 host	IGMPv2 VLAN	Out of range	Ignore the SSM channel table and process the report as if it is in PIM-SM mode.
IGMPv3 host	IGMPv3 VLAN	Out of range	Drop report.

Host	VLAN	SSM range	Action
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic enabled. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and matches an existing SSM channel entry. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and does not match an existing SSM channel entry. Drop report.
In IGMPv3 backward compatibility mode			
IGMPv2 host	IGMPv3 VLAN	In range	Drop report. If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 host	IGMPv3 VLAN	Out of range	Drop report. Process the report as in PIM-SM mode.

When an IGMPv3 group report enters the VLAN port and the IGMP access control must discard one or more of the groups in that packet, the port drops the entire packet and does not forward it on to other ports of the VLAN.

After an IGMPv3 interface receives an IGMPv2 or v1 query, the interface backs down to IGMPv2 or v1. This action flushes all senders and receivers on this interface.

Avaya Ethernet Routing Switch 8800/8600 implementation of PIM-SSM over SMLT

The following sections describe how PIM-SSM and SMLT are implemented in the Avaya Ethernet Routing Switch 8800/8600.

The Avaya Ethernet Routing Switch 8800/8600 SMLT (Split Multi-Link Trunking)/RSMLT (Routed Split Multi-Link Trunking) solution is designed to provide fast fail-over for L2/L3 traffic. It also increases scalability and reliability of a L2 network by providing multiple paths from edge to core, thus eliminating single point of failure. Fast failover for PIM-SSM protocol in Triangle/Square/Full mesh SMLT/RSMLT does not work in PIM-SSM over SMLT for few scenarios. A minimum of 2 to 12 seconds of traffic loss is encountered based on the setup configurations.

To support IP multicast over SMLT/RSMLT, the two aggregate SMLT switches at the core must be visualized as a single switch by the edge (closet) switch. This basic SMLT concept as whole is same and is applicable to IP multicast solution. To achieve this destination, two SMLT aggregate switches enhance the current SMLT protocol and message exchange between these switches to provide a common view on the core switches from a multicast routing perspective. You can achieve fast failover for Multicast traffic in PIM-SSM network through SMLT/RSMLT support for PIM-SSM protocol. PIM-SSM is support on Triangle/Square/Full mesh topologies of SMLT/RSMLT.

PIM-SSM message exchange

The following messages are received whenever an edge switch is flooded on SMLT peer:

Table 4: PIM-SSM message exchange

Messages	Sent to
HELLO	224.0.0.13
JOIN-PRUNE	224.0.0.13
ASSERT	224.0.0.13

The preceding messages are specified to PIM_ALL_ROUTERS (224.0.0.13). Due to flooding of messages on VLAN, the messages are received on both the IST switches.

The following actions are taken when an IST switch receives the PIM-SSM messages on an IST port:

- The packets received from the SMLT edge on an IST port are mapped onto the received SMLT port and the message is processed.
- The ARP entry corresponding to the SRC IP address is checked and the packets are mapped to the SMLT port.
- The source MAC address is acquired on the SMLT port and the PIM message is processed.

JoinPrune, Hello, and Assert messages, which are received on an IST from an SMLT neighbor are processed when they are received on an SMLT port. Similarly, on receiving a JP message on an IST port and if the entry (S,G) is updated, then the SMLT port is considered. Both the IST switches are involved in ASSERT war on SMLT port when received from edge switch.

PIM-SSM static source groups

You can configure static source group entries in the PIM-SSM multicast routing table with static source groups. PIM-SSM cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 25.

Configuration limitations

Avaya recommends that you run PIM-SSM on either all the switches in the domain or only on the edge routers. If you use a mix of PIM-SSM and PIM-SM switches in the domain, run PIM-SSM on all the edge routers and run PIM-SM on all the core routers.

Important:

A PIM domain with edge routers running PIM-SM and core routers running PIM-SSM does not work properly.

Avaya does not support SSM interoperability with DVMRP. However, the MBR functionality works properly for non-SSM groups because SSM-enabled interfaces use PIM-SM behavior for groups outside the SSM range.

SSM switches running IGMPv3 drop reports that they receive out of the SSM range. The SSM switch does not forward them to a PIM-SM switch.

Static source groups cannot conflict with SSM channels and vice versa. After you configure a static source group or an SSM channel, the switch performs a consistency check to make sure no conflicts exist. You cannot map one group (G) to different sources for both a static source group and an SSM channel.

PIM passive interfaces

You can configure the PIM interface as active or passive. The default is active. With an active interface, you can configure transmit and receive PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you use a high number of PIM interfaces and these interfaces connect to end users, not to other switches.

A PIM passive interface does not transmit and drops messages of the following type:

- hello
- join/prune
- register
- register-stop
- assert
- candidate-RP-advertisement
- bootstrap

If a PIM passive interface receives these types of messages, it drops them and the switch logs a message, detailing the type of protocol message that is received and the IP address of the sending device. These log messages help to identify the device that is performing routing on the interface, which is useful if you must disable a device that is not operating correctly.

Important:

A device can send register and register-stop messages to a PIM passive interface, but these messages cannot send out of that interface.

The PIM passive interface maintains information about hosts, through IGMP, that are related to senders and receivers, but the interface does not maintain information about PIM neighbors. You can configure a bootstrap router (BSR) or a rendezvous point (RP) on a PIM passive interface.

You can also use the PIM passive interface feature as a security measure to prevent routing devices from becoming attached and participating in the multicast routing of the network.

Important:

Before you change the state (active or passive) of a PIM interface, disable PIM on that interface. This action prevents instability in the PIM operations, especially when neighbors are present or streams are received. After you disable PIM, the switch loses traffic for approximately 80 seconds.

PIM-SM, PIM-SSM, and DVMRP over SMLT considerations

When deploying PIM-SM, PIM-SSM and DVMRP over an SMLT network configuration, in some instances, the (S,G) record for an mroute may not indicate the correct ingress port. This discrepancy is only a display issue that has no effect on the actual forwarding of traffic over the correct links. Furthermore the display error does correct itself in most cases after a certain time delay. The erroneous ingress port display can occur either on an SMLT link up/down event or on an individual SMLT port up/down event which causes traffic to be switched over to either the IST MLT bundle or to another link within the SMLT bundle. To help determine the correct ingress port for multicast traffic in such a case, look at the traffic statistics of the ports in the SMLT and IST link bundles.

Note that this scenario can occur in all possible SMLT configurations, namely SLT, triangle SMLT, square SMLT and full-mesh SMLT.

Pragmatic General Multicast

Pragmatic General Multicast (PGM) is a standard transport-level protocol that addresses the disadvantages inherent in other multicasting protocols such as unreliable packet delivery, packet duplication, and network congestion. PGM provides reliable, duplicate-free delivery of data packets while reducing network congestion.

PGM sources multicast data packets (ODATA, original data) in an ordered sequence. If a packet is missing in the sequence, the receiver detects the missing sequence number and unicasts a negative acknowledgement (NAK) to the source. The source responds by sending a NAK confirmation (NCF), and either the source or a designated local repairer (DLR) retransmits the missing packet (RDATA, retransmitted data).

PGM is defined by the IETF (RFC 2308).

PGM requires a host implementation as well as a network element (Layer 3 device) implementation.

The Avaya Ethernet Routing Switch 8800/8600 implements the network element aspect of PGM, but does not implement the DLR as it requires a high amount of buffering. End systems usually implement DLR.

PGM concepts and terminology

The following sections describe PGM concepts and how they are used within a PGM network.

Transport session identifiers

PGM runs over IP multicast and delivers data from a source to one or several receivers. The key distinction between PGM and other multicast protocols is that data must deliver within a transmit window time frame. Each multicast session uses a transport session identifier (TSI). PGM supports a number of sources within a multicast group, but each uses its own TSI, and all sources operate independently.

Source path messages

Source path messages (SPM) establish the source path state for the TSI. The source sends out SPMs to maintain up-to-date PGM neighbor information for the distribution tree from source to receivers. PGM receivers also use this information to address negative acknowledgements (NAK) to the source.

Negative acknowledgements

PGM uses NAKs to ensure reliable packet delivery. After a receiver detects a missing packet, it repeats this NAK until it receives a NAK confirmation. By doing this, PGM guarantees that receivers can either receive all data packets from transmissions and retransmissions, or detect unrecoverable data packet loss.

PGM also uses NAKs to reduce congestion. Instead of sending a positive acknowledgement (ACK) every time a packet is received, which adds to network overhead, PGM issues a NAK only after it does not receive a packet.

NAK confirmations

NAK confirmations (NCF) further reduce congestion by suppressing redundant NAKs. After a receiver detects a missing packet, it unicasts a NAK to the next-hop upstream PGM router.

This router multicasts an NCF to the subnet so other receivers do not send additional NAKs. The router also stores the address of the group so that it forwards retransmissions only to those segments containing receivers that require the packet.

Designated local repairers

Designated local repairers (DLR) are local hosts that retransmit missing data packets for a number of multicast groups. Designated local repairers multicast the missing data to receivers below it in the distribution tree. This technique reduces the load on the network due to retransmissions and reduces the time for receivers to recover missing packets.

Important:

The Avaya Ethernet Routing Switch 8800/8600 cannot serve as a DLR because DLRs require a large amount of buffering. Therefore, the null negative acknowledgement (NNAK) parameters in Enterprise Device Manager and the CLI are not supported.

PGM network element

You can configure an Avaya Ethernet Routing Switch 8800/8600 as the network element of a PGM network. The network element performs the following PGM tasks:

- Sources periodically interleave SPMs with data frames. Source path message frames allow the network element (an Avayal Ethernet Routing Switch 8800/8600 with PGM enabled) to learn the path to the source and to maintain information about the PGM session.
- Because data frames are numbered, hosts can detect missing data and issue NAKs if data is missing. The network element forwards NAKs to the source and stores information about the NAK to forward retransmitted data.
- Hosts continue to send NAK messages until they receive a NCF from the network element. To reduce traffic on the network, the network element does not forward all received NAKs because they transmit from several receivers for the same lost packet.
- The source retransmits the requested frame that is forwarded by the network element to the interfaces that send the NAK.

Multicast VLAN Registration Protocol

The Avaya Ethernet Routing Switch 8800/8600 uses IGMP Snoop to listen for report, leave and query packets, and then creates or deletes multicast groups for receiver ports to receive multicast data streams. In IGMP Snoop, all the ports, including receiver and source ports, are members of the same VLAN. When users in different VLANs join the same group, the multicast

router replicates one stream into multiple streams that are sent to these VLANs. Multiple streams waste bandwidth and decrease the performance of the multicast router.

Important:

If IGMPv3 snoop interface is used for MVR both the MVR VLAN and the Receiver VLANs must be set to version 3.

The Multicast VLAN Registration (MVR) Protocol solves this problem. With MVR, the receiver ports remain in the IGMP Snoop VLAN, but one VLAN is designated as the MVR VLAN. Each Ethernet Routing Switch 8800/8600 supports only one MVR VLAN.

The MVR VLAN has a source port, which connects to the multicast router. After you bind several IGMP Snoop VLANs to the MVR VLAN, and a multicast data packet arrives from the source port, the switch replicates this packet and forwards it to all the IGMP Snoop VLANs that are bound to the MVR VLAN.

For example, if two users in two separate VLANs join the same group, IGMP Snoop creates two groups with one group for each VLAN. The multicast router then connects to the MVR VLAN in the switch and not to the two VLANs. Consequently, multicast sends only one copy of the multicast data stream to the switch and the switch replicates this data stream in the MVR VLAN. The MVR VLAN sends the data stream to the two groups where the users receive it.

MVR is based on IGMP, but the two features work independently. The MVR VLAN controls only the VLANs that are bound to it; other IGMP Snoop VLANS operate as usual.

After you enable MVR globally, all IGMP control packages that are received from IGMP Snoop VLANs bound to the MVR VLAN (including report, leave, and query) are processed by MVR.

MVR and multicast routing such as Protocol Independent Multicast (PIM) are exclusive. If you enable the MVR Protocol, you cannot enable multicast routing and conversely, if you enable multicast routing, you cannot enable MVR.

Important:

If IGMPv3 snoop interface is used for MVR both the MVR VLAN and the Receiver VLANs must be set to version 3, otherwise it may not work correctly.

MVR and IGMP fast leave on Mrouter ports

When an IGMP interface receives an IGMP Leave message, the normal behavior is for the interface to send a group-specific query and wait for any reports before stopping traffic forwarding for the group. When IGMP fast leave is enabled on the interface, this normal IGMP behavior is skipped. Fast leave mode allows a switch to immediately stop forwarding a specific group's multicast traffic as soon as an IGMPv2 Leave Group message is received on an interface. The interface is immediately removed from the list of receivers for that group.

This assumes that there is only one receiver for this group on this interface or that all receivers agree to stop receiving traffic for the group when one of them leaves the group.

Fast leave can be configured at the port level or the VLAN level (in which case all the ports in the VLAN must be fast-leave compatible).

Fast leave is applicable to those receivers that are directly attached to a port or VLAN on which fast-leave is enabled. That is, if any IGMP receiver is attached to any receiver VLAN, fast-leave must be enabled on this port or VLAN.

Due to this behavior, Avaya recommends to disable fast-leave for Mrouter ports connected to an MVR VLAN.

Limitations to MVR

The following restrictions and limitations apply to MVR:

- The MVR feature is designed as an edge service only. Do not enable MVR in the core network.
- The maximum number of multicast streams shared by MVR, IGMP snooping and Layer 3 Multicast protocols is 2000 for every switch in SMLT environments and 4000 for every switch in non-SMLT environments.
- The maximum number of receiver VLANs in a VRF is 64.
- The maximum number of MVR VLANs in a VRF is 1.
- The total maximum number of receiver VLANs in all 256 VRFs is 2048.
- The maximum number of joins per sec is 250 for every switch.
- When IGMP V3 is enabled and because IGMP only stores one source for every group in the SSM list, the local system sends group-and-source-specific query which only contain one source for every group to return this sources IGMP report (its record type is BLOCK_OLD_SOURCES).
- When IGMPv3 is enabled, IGMP only supports report messages with record type of MOD_IS_INCLUDE, ALLOW_NEW_SOURCES or BLOCK_OLD_SOURCES, so the IGMP querier only processes these messages.
- When the MVR VLAN is configured for IGMPv3, IGMPv1 receivers are not accepted. The IGMP version on the receiver VLAN must be no higher than the IGMP version on the MVR VLAN.
- When MVR proxy is enabled or disabled, fast leave must be disabled on the Mrouter port connected to the MVR VLAN.
- The MVR VLAN requires an IP address, which must not be deleted. The MVR feature is not functional if the MVR VLAN has an IP address of 0.0.0.0.
- IGMP snooping is automatically enabled on the MVR VLAN, and it cannot be disabled. In addition, snooping must be enabled on a VLAN to add it as receiver VLAN in the MVR. Snooping can only be disabled after MVR is disabled or when the VLAN is no longer a receiver VLAN.
- IGMP querier must not be enabled on the MVR VLAN.

Multicast MAC filtering

Some network applications rely on a Layer 2 multicast MAC mechanism to send a frame to multiple hosts for processing. For example, mirroring is one such application. With Release 3.3 or later of the Avaya Ethernet Routing Switch 8800/8600 software, you can direct MAC multicast flooding to a specific set of ports using the multicast MAC filtering feature.

Important:

You can configure multicast MAC filtering only for local addresses on a switch. You cannot use this feature to route traffic between switches (for example, you cannot configure it to forward for interfaces that are not local).

The multicast MAC is configured as a MAC address where the least significant bit of the most significant byte is set to 1. The multicast MAC filtering feature is available for Layer 2. Because it is also effective for IP routed traffic, however, Layer 3 functionality is available as well. This filtering does not apply to Bridge Protocol Data Units (BPDU).

In Layer 2, a multicast MAC address generally floods to all ports in the VLAN. With multicast MAC filtering, you can define a separate flooding domain for a multicast MAC address, which is a subset of the ports on a VLAN. The maximum number of multicast MAC addresses that you can configure is 100. However, depending upon the overall configuration of your switch, you can be limited to fewer addresses.

In Layer 3, you must configure an Address Resolution Protocol (ARP) entry for routed traffic that maps the unicast IP address to the multicast MAC address and lists the delivery ports for data destined for that IP or multicast MAC address.

To perform multicast MAC filtering, create the VLAN, and then manually define a flooding domain (that is, MAC address and port list) for a specific multicast address. When you specify the multicast MAC flooding domain, you must indicate the ports or multilink trunks to consider for multicast traffic. The flooding is based on whether the specified ports are active members in the VLAN.

High Availability

Beginning with Release 5.0, the Avaya Ethernet Routing Switch 8800/8600 provides partial High Availability (HA) for IP multicast. In this mode, the applications perform synchronization of configuration information only; no dynamic information is synchronized. This mode forces applications to restart on failover.

Release 5.0 supports partial HA for DVMRP, PIM-SM, PIM-SSM, and mroute commands. Release 5.1 adds partial HA for multicast virtualization.

Virtualization

Beginnning with Release 5.1, you can configure multicast routing support with the Virtual Router Forwarding (VRF) Lite feature. You can use VRF Lite to emulate many virtual routers with one router; VRF Lite virtualizes the following multicast routing protocols: IGMP, PIM-SM, and PIM-SSM.

The Ethernet Routing Switch 8800/8600 supports multicast virtualization with the following hardware and software requirements:

- Multicast virtualization works on an R, RS, 8800 module port.
- Multicast virtualization requires an 8692 SF/CPU with SuperMezz or an 8805 SF/CPU.
- You must configure virtualization for the unicast routing system.

If you configure multicast virtualization, many IGMP and PIM commands that you use on the Global Router automatically apply to VRFs that you create. Occasionally, you must also configure the multicast protocol on the VRF instance; the procedures in this document identify when you must use a multicast command on a VRF instance to effect the virtualized configuration.

To configure multicast virtualization, you must purchase and install the Premier License. For more information about how to obtain and install licenses, see *Avaya Ethernet Routing Switch 8800/8600 Administration, NN46205-605*.

Limitations

The following restrictions and limitations apply to multicast virtualization:

- Release 5.1 does not virtualize DVMRP, PIM-MBR, IGAP, and PGM.
- Avaya supports 200 active PIM interfaces and 200 active PIM neighbors for all VRFs.
- Avaya supports 2000 passive PIM interfaces for all VRFs.

Virtualization scenarios

This section illustrates three scenarios in which you can use multicast virtualization. Although the figures in this section depict the routers as provider edge (PE) routers, this is not a requirement.

Basic domain configuration

You can create separate multicast domains over a shared infrastructure. In Release 5.1, you can create only PIM-SM and PIM-SSM domains in different virtual networks. All the protocol messages and data traffic are by default kept within a virtualized multicast domain. PIM-SM

and PIM-SSM depend on unicast routing information. You must use the unicast routing protocols on the VRF. The following figure illustrates this configuration.

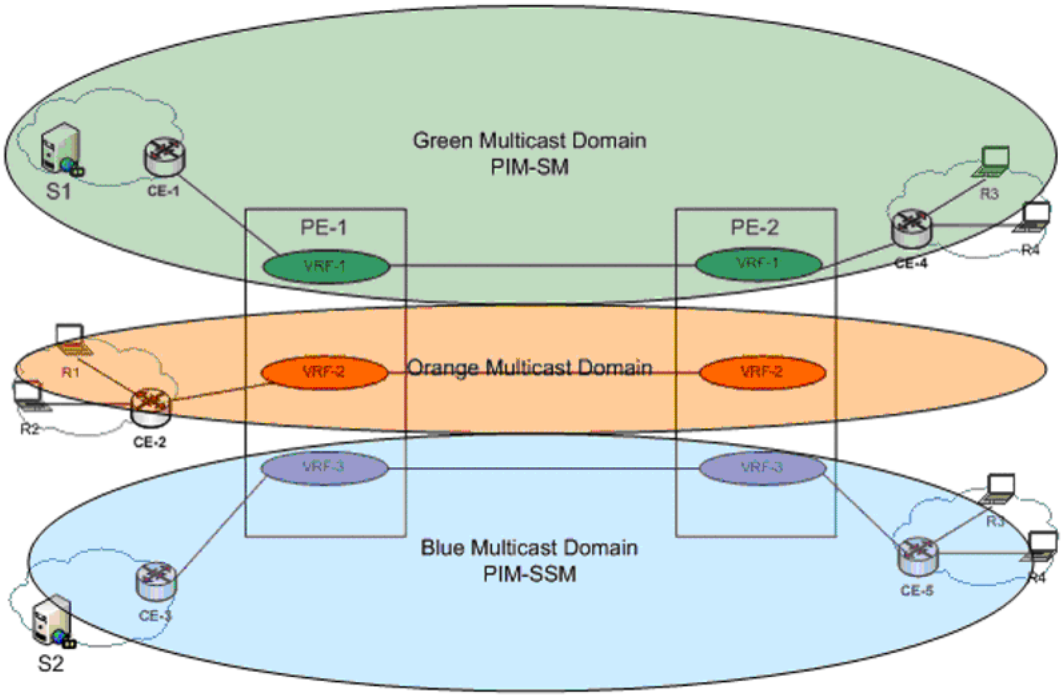

Figure 15: Virtualized multicast domains

Interdomain multicast in virtualized networks

The Avaya Ethernet Routing Switch 8800/8600 supports interdomain multicast routing or connecting domains that run different multicast protocols, only in the Global Router, VRF 0. This limitation is because only PIM-SM or PIM-SSM domains can connect to a VRF in a PE. You can overcome this limitation by connecting a customer edge (CE) device with the MBR functionality.

Use PIM-SM on the interface that connects the CE device to the VRF in PE. Use DVMRP in the link that connects to the device in the DVMRP or PIM-DM domain. Use the MBR functionality in the CE device. The following figure illustrates this configuration.

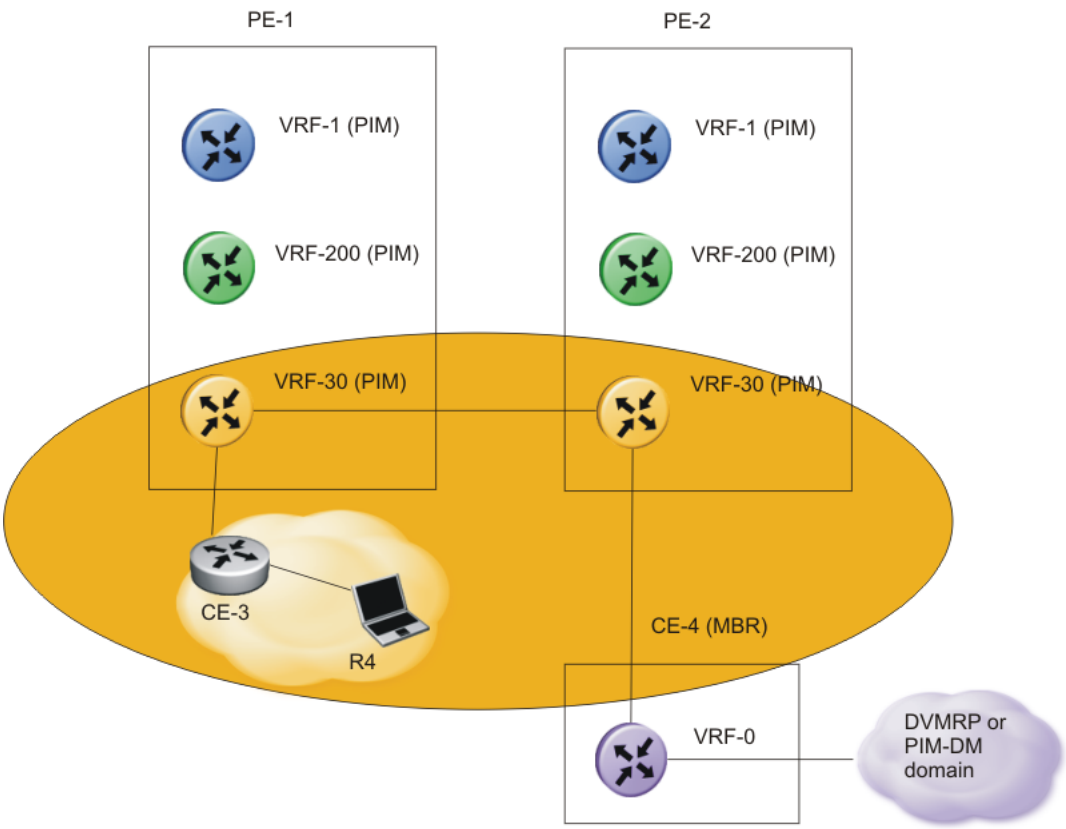

Figure 16: Interdomain multicast in virtualized networks

SMLT

The Avaya Ethernet Routing Switch 8800/8600 provides IP multicast routing support on SMLT by building a virtual switch that represents the two switches of the SMLT core. The virtual switch maintains identical multicast routing protocol states on both core switches connected together by the IST link.

Every PIM or PIM-SSM instance in every VRF must keep the routing and adjacency states in both core switches identical. PIM sends the PIM-Hello, PIM-JP, and PIM-ASSERT messages to the ALL-PIM-ROUTER multicast address, and because PIM sends the message on the SMLT VLAN, both core switches receive it. The switch that receives these messages on the IST must associate the correct SMLT port based on the source MAC address.

Each VRF simulates the core switch by synchronizing their multicast routing protocols states over the IST link. The following figure illustrates this configuration.

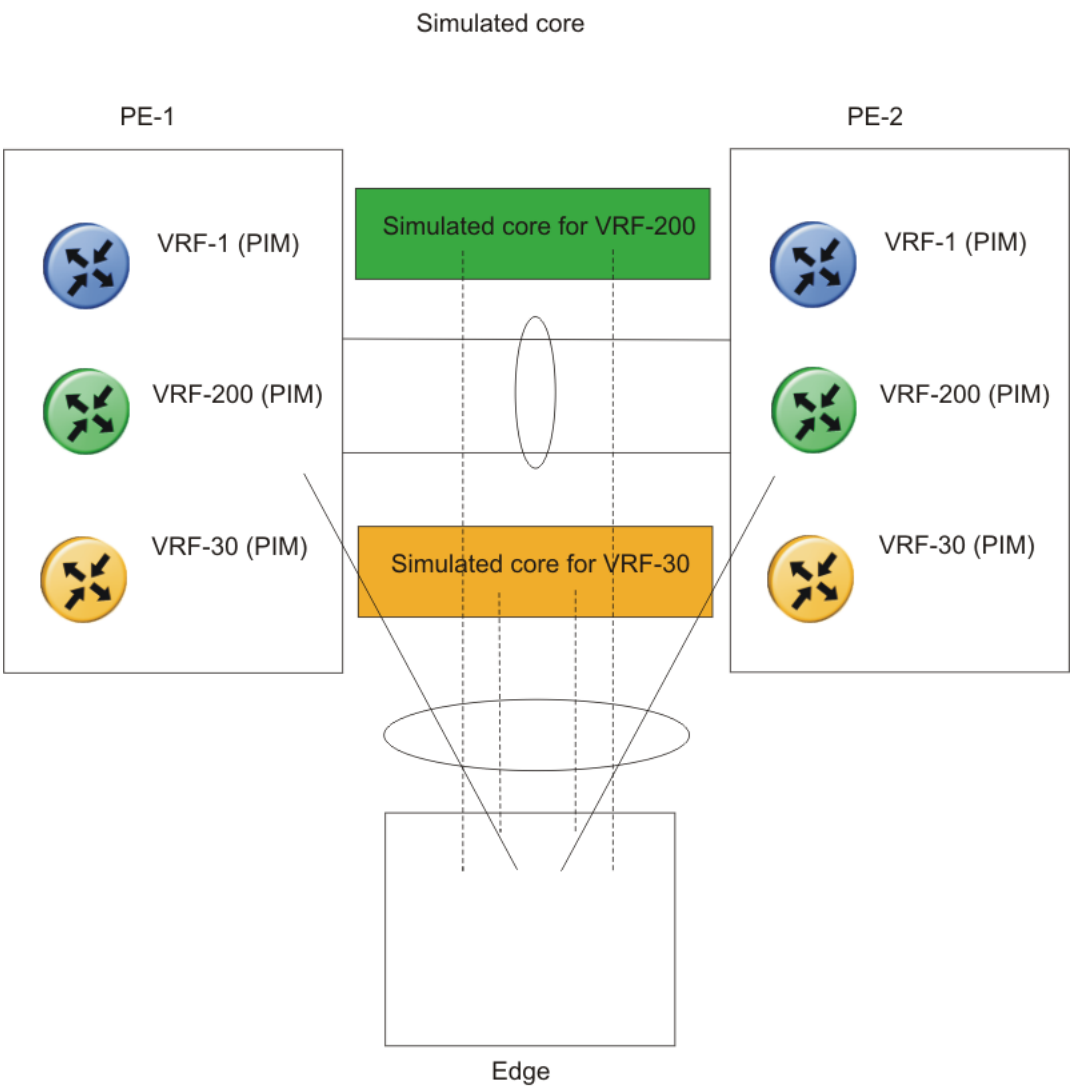

Figure 17: SMLT in virtualized multicast

IGMP L2 Querier

IGMP snoop (Internet Group Management Protocol snooping) enables Layer 2 switches in the network to snoop into IGMP control protocol packets that are exchanged between downstream hosts and upstream routers. By snooping packets, the L2 switches generate the L2 MAC forwarding table that are used for further switching sessions and regulate the multicast traffic from flooding the Layer 2 segment of the network. The multicast traffic fails only on downstream ports where listeners are present.

Cyclic query is required for maintaining a group in the IGMP protocol, which is dealt by the multicast router. However, the layer 2 switch does not send out the query automatically if the multicast router is not connected. This prevents the system from participating in many multicast

customer environment where users do not require any multicast routing. The query is supported only when the Avaya Ethernet Routing switch 8800/8600 works on layer 2, since IGMP L2 Querier induces a multicast router on layer 2.

IGMP IGMP L2 Querier limitations

The IGMP L2 Querier limitations are as follows:

- After the snoop feature and querier feature are enabled on an interface and if the system does not receive any IGMP query message, the system becomes the querier. When one IGMP query message is sent to an interface and an another querier is sent to multicast VLAN, and the other Querier present timer present in the multicast VLAN timer is not expires, then the other Querier present timer acts as a nonquerier until the other Querier present timer expires.
- Currently the ERS 8800/8600 does not support Querier selection, therefore Avaya recommends that, you must enable Querier selection on only one switch in the L2 multicast domain.
- After you enable IGMP V3, IGMP stores one source for each group in the SSM list. The local system then sends group-and-source-specific query that contains one source for each group to return the source IGMP report. It’s record type is BLOCK_OLD_SOURCES).
- IGMP querier only processes message types such as MOD_IS_INCLUDE , ALLOW_NEW_SOURCES or BLOCK_OLD_SOURCES.
- IGMP L2 querier feature is based on IGMP SNOOP. If the IGMP SNOOP is disabled, the L2 Querier does not work until the IGMP SNOOP and IGMP L2 Querier are enabled again.

Chapter 4: IP multicast routing configuration

Configure IP multicast routing to transmit data from a source to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting; however, multicasting transmits data to specific groups, and broadcasting transmits to all devices on the network. Because multicasting transmits only one stream of data to many destinations, multicasting conserves bandwidth. This section outlines the required and optional tasks you need to perform to configure IP multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

Prerequisites to IP multicast routing configuration

- You must configure at least one IP interface on the Avaya Ethernet Routing Switch 8800/8600. For more information about configuring interfaces, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing,* (NN46205-523).

IP multicast routing configuration tasks

This work flow shows you the sequence of tasks you perform to configure IP multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

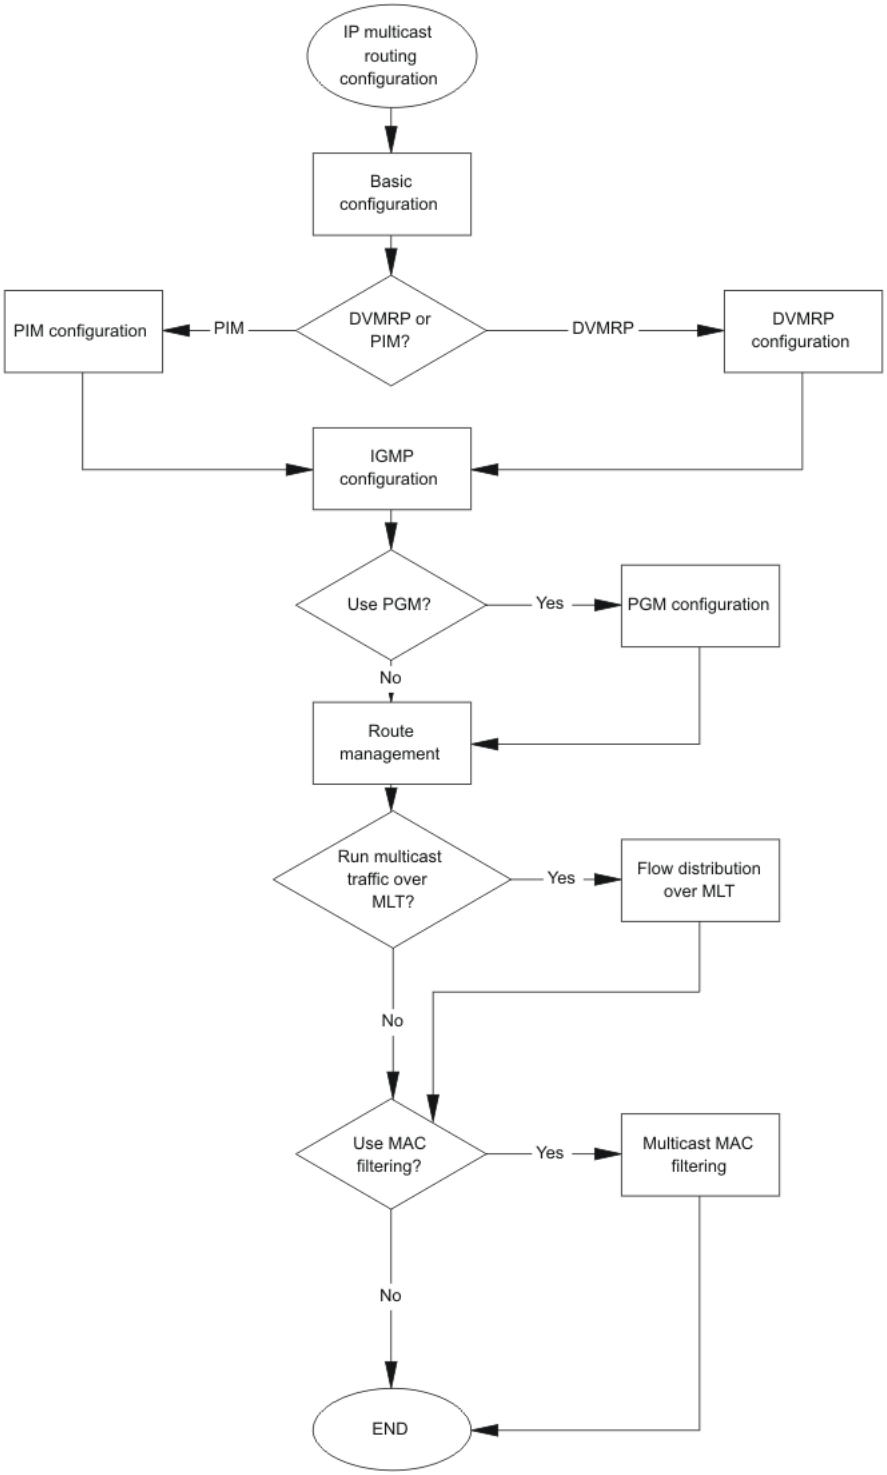

Figure 18: IP multicast routing configuration tasks

IP multicast routing configuration navigation

- IP multicast basic configuration using Enterprise Device Manager on page 87
- IP multicast basic configuration using the CLI on page 107
- IP multicast basic configuration using the ACLI on page 145
- DVMRP configuration using Enterprise Device Manager on page 179
- DVMRP configuration using the CLI on page 201
- DVMRP configuration using the ACLI on page 231
- PIM configuration using Enterprise Device Manager on page 253
- PIM configuration using the CLI on page 267
- PIM configuration using the ACLI on page 281
- IGMP configuration using Enterprise Device Manager on page 291
- IGMP configuration using the CLI on page 315
- IGMP configuration using the ACLI on page 355
- PGM configuration using Enterprise Device Manager on page 379
- PGM configuration using the CLI on page 387
- PGM configuration using the ACLI on page 395
- Route management using Enterprise Device Manager on page 401
- Route management using the CLI on page 413
- Route management using the ACLI on page 427
- Multicast flow distribution over MLT using Enterprise Device Manager on page 439
- Multicast flow distribution over MLT using the CLI on page 445
- Multicast flow distribution over MLT using the ACLI on page 451
- Multicast MAC filtering using Enterprise Device Manager on page 479
- Multicast MAC filtering using the CLI on page 483
- Multicast MAC filtering using the ACLI on page 489

Chapter 5: IP multicast basic configuration using Enterprise Device Manager

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts subscribe to multicast services using a host membership protocol. The Internet Group Management Protocol (IGMP) is an example of a host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast-Sparse Mode (PIM-SM) and Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

IP multicast basic configuration using Enterprise Device Manager procedures

The following task flow shows you the sequence of procedures you perform to configure basic elements of IP multicast routing.

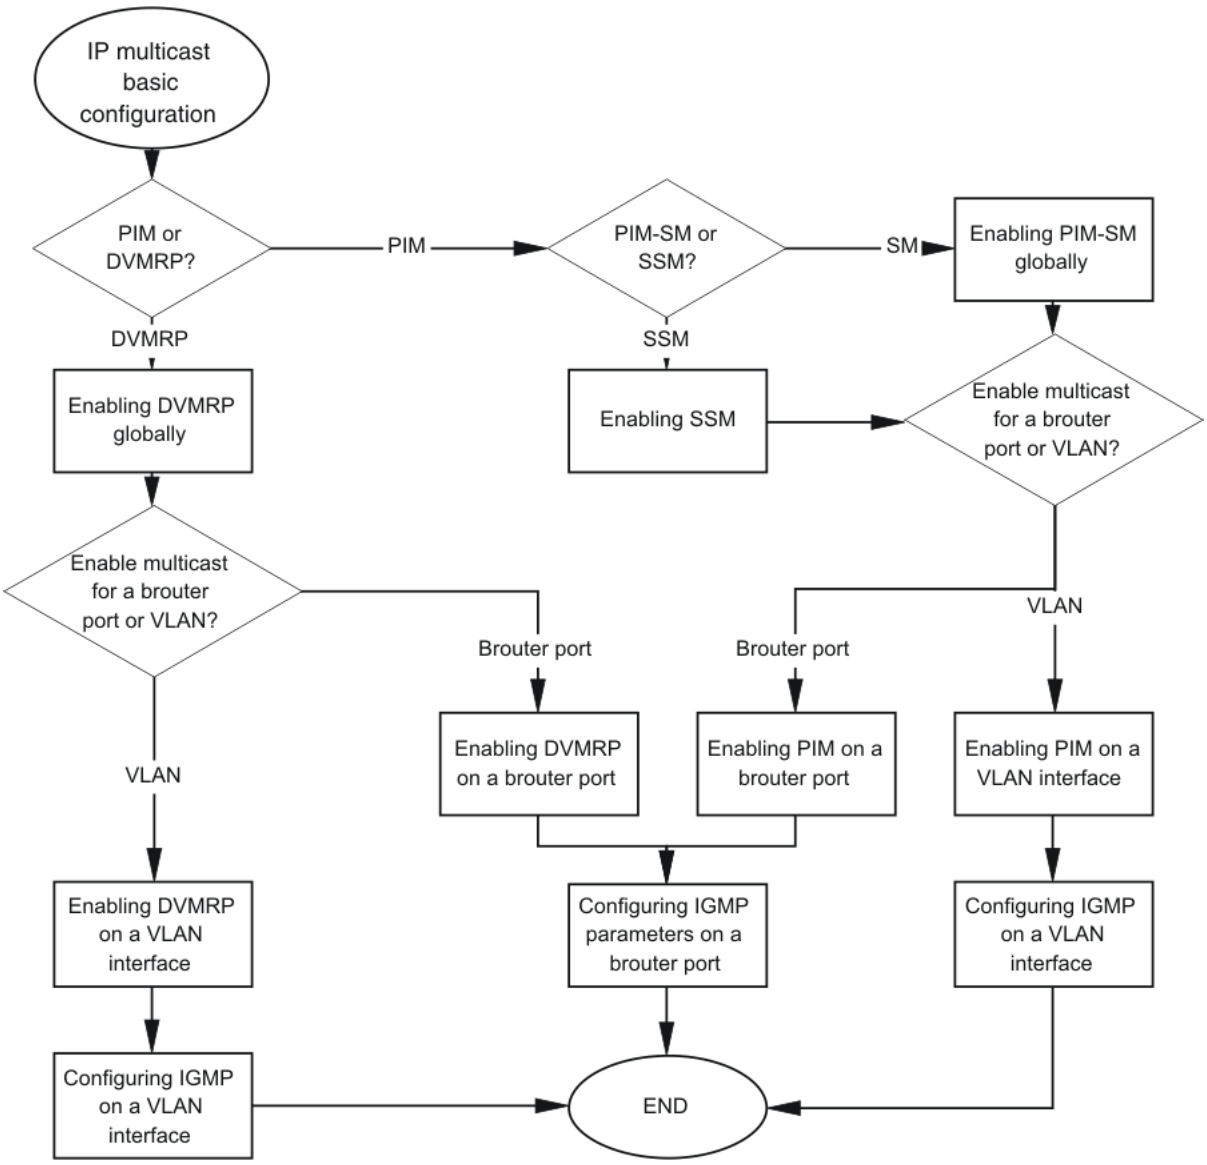

Figure 19: IP multicast basic configuration using Enterprise Device Manager procedures

Navigation

- Enabling DVMRP globally on page 89
- Enabling PIM-SM globally on page 90
- Enabling PIM on a brouter port on page 93
- Enabling SSM globally on page 94

- Enabling PIM on a VLAN interface on page 95
- Enabling DVMRP on a brouter port on page 96
- Enabling DVMRP on a VLAN on page 98
- Configuring IGMP parameters on a brouter port on page 99
- Configuring IGMP parameters on a VLAN on page 102

Enabling DVMRP globally

Configure the Distance Vector Multicast Routing Protocol (DVMRP) globally to offer multicasting services on all interfaces.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Select **Enable**.
4. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Globals tab.

Variable	Value
Enable	Enables (true) or disables (false) DVMRP on the routing switch. You must globally enable DVMRP before you can enable a port or a VLAN as IGMP or DVMRP.
UpdateInterval	Periodically, each multicast router advertises routing information on each DVMRP interface, using the DVMRP export message. You can specify the time interval (in seconds) between DMVRP updates. The range is from 10–2000 with a default of 60.
TriggeredUpdateInterval	The switch sends triggered updates after routing information changes. This value is the amount of time (in seconds) between triggered update messages. The range is from 5–1000 with a default of 5.
LeafTimeOut	After DVMRP advertises a route on an interface, it waits a period of time for a DVMRP neighbor to respond positively.

Variable	Value
	If no neighbor responds in the allowed time, the router considers the network attached to the interface as a leaf network. Use the leaf timer to specify how long (in seconds) the router waits for a response from a neighbor. The range is from 25–4000 with a default of 200.
NbrTimeOut	The neighbor report timer specifies how long (in seconds) the router waits to receive a report from a neighbor before considering the connection inactive. The range is from 35–8000 with a default of 140.
NbrProbeInterval	Specifies how often the DVMRP router sends probe messages on its interfaces. The range is from 5–30 seconds with a default of 10.
RouteExpireTimeOut	Defines the route expiration timeout value. The range is from 20–4000 seconds with a default of 140.
FwdCacheTimeOut	Defines the forward cache timeout value, which is used to age prune entries. The range is from 300–86400 seconds with a default of 300.
RouteDiscardTimeOut	Defines the time to garbage collect route. The range is from 40–8000 seconds with a default of 260.
RouteSwitchTimeOut	Defines the route discard timeout value. The range is from 20–2000 seconds with a default of 140.
VersionString	The DVMRP version information for the router.
GenerationId	Used by neighboring routers to detect whether a reset or disable or enable DVMRP action occurred to the switch or to a particular interface. If so, the router must resend the entire multicast routing table to its neighbor immediately instead of waiting for the next scheduled update.
NumRoutes	The number of entries in the routing table. Use this information to monitor the routing table size to detect illegal advertisements of multicast routes.
ReachableRoutes	The number of entries in the routing table with noninfinite metrics. Use this number to detect network partitions by observing the ratio of reachable routes to total routes.

Enabling PIM-SM globally

Enable PIM-SM to offer multicasting services. After you enable PIM-SM globally and on a particular interface, the IGMP parameters take effect.

Prerequisites

- To configure PIM-SM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. In the Mode field, click **sm** (sparse mode).

   Important:

   You can use static RP when SSM is enabled for groups outside the SSM range.
4. Select the **Enable** check box.
5. Select the **Mbr (Multicast Border Router)** check box.
6. In the **JoinPruneInterval** field, enter how long to wait in seconds before the PIM router sends out the next join/prune message to its upstream neighbors or use the default value.
7. In the **RegisterSuppTimer** field specify how long in seconds the designated router (DR) suppresses sending registers to the RP or use the default value.
8. In the **UniRouteChgTimeOut** field, enter how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates for PIM or use the default value.
9. In the **DiscardDataTimeOut** field specify how long in seconds to discard data until the join is received from the RP or use the default value.
10. In the **CRPADVTimeOut** field, enter how often (in seconds) a router configured as a candidate rendezvous point (C-RP) router sends C-RP advertisement messages, or use the default value.
11. In the **BootStrapPeriod** field, enter an interval in seconds that the elected bootstrap router (BSR) waits between originating bootstrap messages, or use the default value.
12. Select **StaticRP** check box, to enable static RP feature.
13. In the **FwdCacheTimeOut** field, enter the PIM forward cache expiry value in seconds or use the default value.

14. In the **FastJoinPrune** field, enable or disable the PIM fast join prune feature.
15. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Globals tab.

Variable	Value
Mode	Configures the mode on the routing switch: sm (Sparse Mode) or ssm (Source Specific Multicast).
Enable	Enables or disables PIM.
Mbr (Multicast Border Router)	Configures the router as a PIM multicast border router (MBR). PIM MBRs connect PIM domains to other multicast routing domains and the rest of the Internet. In particular, the MBR on an Avaya Ethernet Routing Switch 8800/8600 permits the connection of a PIM-SM domain to a Distance Vector Multicast Routing Protocol (DVMRP) domain.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join/prune message to its upstream neighbors. The range is from 1–18724 and the default is 60 seconds.
RegisterSuppTimer	Specifies how long (in seconds) the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. The range is from 6–65535 and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates for PIM. The range is from 2–65535 and the default is 5 seconds. **Important:** If you lower this value, it increases how often the switch polls the RTM. This value can affect the performance of the switch, especially when a high volume of traffic flows through the switch.
DiscardDataTimeOut	Specifies how long (in seconds) to discard data until the join is received from the RP. An IP multicast discard record is created after a register packet is sent, until the timer expires or a join is received. The range is from 5–65535 and the default is 60 seconds.
CRPADVTimeOut	Specifies how often (in seconds) a router configured as a candidate rendezvous point (C-RP) router sends C-RP

Variable	Value
	advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected BSR. The range is from 5–26214 and the default is 60 seconds.
BootStrapPeriod	Specifies the interval (in seconds) that the elected bootstrap router (BSR) waits between originating bootstrap messages. The range is from 5–32757 and the default is 60 seconds.
StaticRP	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism. The default is disabled (clear).
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value is used in aging PIM mroutes in seconds. The range is from 10–86400 and the default value is 210.
FastJoinPrune	Enables or disables the PIM fast join prune feature. The default is disable.

Enabling PIM on a brouter port

Configure PIM for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable PIM globally before you enable it on an interface.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Select the **PIM** tab.
5. Select **Enable** check box.
6. Click **Apply**.

Variable definitions

Use the data in the following table to configure the PIM tab.

Variable	Value
Enable	Enables or disables PIM.
Mode	Specifies the mode on the interface: sparse or source specific multicast (SSM).
IntfType	Specifies the PIM interface type: active or passive.
HelloInterval	The frequency at which PIM Hello messages are transmitted on this interface.
JoinPruneInterval	The frequency at which PIM Join/Prune messages are transmitted on this PIM interface.
CBSRPreference	The preference value for the local interface as a candidate bootstrap router. The value of -1 is used to indicate that the local interface is not a candidate BSR interface.

Enabling SSM globally

Enable Source Specific Multicast (SSM) to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Prerequisites

- Configure a unicast protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) globally and on the interfaces where you want to configure PIM.

For more information about RIP and OSPF, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — OSPF and RIP,* (NN46205-522).

- Enable PIM globally.
- To configure PIM-SM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. In the **Mode** field, select **ssm** (source specific multicast).
4. Select the **Enable** check box.
5. Click **Apply**.

Important:

After you enable PIM in SSM mode, the IGMP parameters take effect. To take full advantage of SSM, enable IGMPv3 if hosts attached to the switch run IGMPv3 or configure the SSM table.

Enabling PIM on a VLAN interface

Configure PIM for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable PIM globally before you enable it on an interface.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select the VLAN that you want to configure with PIM.
4. Click **IP** from the menu bar.

5. Click the **PIM** tab.
6. Select the **Enable** check box.
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the VLAN PIM tab.

Variable	Value
Enable	Enables (true) or disables (false) PIM.
Mode	Displays the mode currently running on the routing switch. The valid modes are SSM and Sparse. This field is a read-only field.
IntfType	Specifies the type of interface: active or passive.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0–18724.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join/prune message to its upstream neighbors. The default is 60 seconds. The range is 1–18724.
CBSRPreference	Configures your preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. The range is 1–255.

Enabling DVMRP on a brouter port

Configure DVMRP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on an interface.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DVMRP** tab.
5. Select the **Enable** check box to select DVMRP on the port, or clear the check box.
6. Enter a metric (cost) of the maximum number of hops for DVMRP; the range is from 1 to 31.

   The default value of 1 means local delivery only. You can use the metric value to control the scope of the DVMRP routes.
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Port DVMRP tab.

Variable	Value
LocalAddress	Provides the IP address of the DVMRP router interface.
Enable	Enables (check box selected) or disables (check box cleared) DRMRP on the port.
Metric	Specifies the distance metric for this port; used to calculate distance vectors. The range is 1–31 hops. The default is 1.
InterfaceType	Configures the port type as passive or active.
DefaultListen	Configures the port to listen (check box selected) or not listen (check box cleared) for the default route. The default is listen.
DefaultSupply	Configures the port to supply (check box selected) or not supply (check box cleared) only the default route. The default is not supply.

Variable	Value
DefaultRouteMetric	Configures the metric (number of hops for DVMRP) of the default route. The range is 1–31 hops. The default is 1.
AdvertiseSelf	Configures the port to advertise (check box selected) or not advertise (check box cleared) local routes to neighbors. The default is advertise.
InPolicy	Specifies the name of the DVMRP accept policy applied to the port.
OutPolicy	Specifies the name of the DVMRP announce policy applied to the port.

Enabling DVMRP on a VLAN

Configure DVMRP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on an interface.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP** from the menu bar.
5. Click the **DVMRP** tab.
6. Select the **Enable** check box to select DVMRP on the port, or clear the check box.
7. Enter a metric (cost) of the maximum number of hops for DVMRP; the range is from 1 to 31.

   The default value of 1 means local delivery only. You can use the metric value to control the scope of the DVMRP routes.
8. Click **Apply**.

Variable definitions

Use the data in the following table to configure the VLAN DVMRP tab.

Variable	Value
LocalAddress	Provides the IP address of the DVMRP router interface.
Enable	Enables (check box selected) or disables (check box cleared) DRMRP on the VLAN.
Metric	Specifies the distance metric for this VLAN; used to calculate distance vectors. The range is from 1–31 hops. The default is 1.
InterfaceType	Configures the VLAN type as passive or active. The default is active.
DefaultListen	Configures the VLAN to listen (check box selected) or not listen (check box cleared) for the default route. The default is listen.
DefaultSupply	Configures the VLAN to supply (check box selected) or not supply (check box cleared) only the default route. The default is not supply.
DefaultRouteMetric	Configures the metric (number of hops for DVMRP) of the default route. The range is from 1–31 hops. The default is 1.
AdvertiseSelf	Configures the VLAN to advertise (check box selected) or not advertise (check box cleared) local routes to neighbors. The default is advertise.
InPolicy	Specifies the name of the DVMRP accept policy applied to the VLAN.
OutPolicy	Specifies the name of the DVMRP announce policy applied to the VLAN.

Configuring IGMP parameters on a brouter port

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- For IGMP parameters to take effect, enable either DVMRP or PIM on the interface.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **IGMP** tab.
5. In the **QueryInterval** field, enter frequency (in seconds) at which the interface transmits IGMP host query packets.
6. In the **QueryMaxResponseTime** field, enter the maximum response time (in tenths of a second) advertised in IGMPv2 general queries.
7. In the **Robustness** field, enter a parameter to tune for the expected packet loss of a network.
8. In the **LastMembQueryIntvl** field, enter the maximum response time (in 1/10 seconds) that must be inserted into group-specific queries sent in response to leave group messages.
9. In the **Version** field, enter the version of IGMP (1, 2 or 3) that you want to configure on this interface.
10. Select the **FastLeaveEnable** check box to enable fast leave.
11. Click enable in **StreamLimitEnable** to enable stream limit.
12. In **Maximum Number Of Stream** box, type the maximum number of streams you want to include on a brouter port.
13. Select **DynamicDowngradeEnable** check box to enable dynamic downgrade of the IGMP version.
14. Select **CompatibilityModeEnable** check box to enable or disable v2-v3 compatibility mode.
15. Click **Apply**.

Variable definitions

Use the data in the following table to configure the IGMP tab.

Variable	Value
QueryInterval	The frequency (in seconds) at which the interface transmits IGMP host query packets. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	The maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the QueryInterval.
Robustness	Use this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If a network is expected to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that one query for each query interval is dropped without the querier aging out.
LastMembQueryIntvl	The maximum response time (in 1/10 seconds) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Avaya recommends that you configure this parameter to values greater than 3. If a fast leave process is not required, Avaya recommends that you use values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
Version	The version of IGMP (1, 2 or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this port.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this port. The range is from 0–65535 and the default is 4.

Variable	Value
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.
DynamicDowngradeEnable	Configures if the Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected, which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear, which means IGMPv3 is not compatible with IGMPv2.

Configuring IGMP parameters on a VLAN

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- For IGMP parameters to take effect, enable either DVMRP or PIM-SM on the interface.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click the **IP** from the menu bar.
5. Click the **IGMP** tab.
6. In the **QueryInterval** field, enter the frequency (in seconds) at which the interface transmits IGMP host query packets.
7. In the **QueryMaxResponseTime** field, enter the maximum response time (in tenths of a second) advertised in IGMPv2 general queries.

8. In the **Robustness** field, enter a parameter to tune for the expected packet loss of a network.
9. In the **LastMembQueryIntvl** field, enter the maximum response time (in 1/10 seconds) that must be inserted into group-specific queries sent in response to leave group messages.
10. Select **SnoopEnable** check box to enable IGMP snooping.
11. Select **SsmSnoopEnable** check box to enable Source Specific Multicast mode.
12. Select **ProxySnoopEnable** check box to enable proxy snoop.
13. Select in the **IgapEnable** check box to enable the Internet Group membership Authentication Protocol (IGAP) on an interface.
14. Select enable in the **AccntEnable** check box to enable the Internet Group membership Authentication Protocol (IGAP) on an interface.
15. Select enable in the **AuthEnable** check box to enable IGAP authentication on an interface.
16. In the **Version** field, enter the version of IGMP (1, 2, or 3) that you want to configure on this interface.
17. Select the **FastLeaveEnable** check box to enable fast leave on the interface.
18. Select the **StreamLimitEnable** check box to enable fast leave on the interface.
19. In the **Maximum Number of Stream** dialog box, type the number of streams.
20. In **Current Number of Stream**, type the number of stream.
21. In the **FastLeavePortMembers** field, click on **[...]** button.

    The Port Editor: FastLeavePortMembers screen appears.
22. Select the ports that you want to enable for fast leave.
23. Click **Ok**. The ports that are enabled appear in the FastLeavePortMembers field.
24. In the **SnoopMRouterPorts** field, click **[...]** button.

    The Port Editor: SnoopMRouterPorts screen appears.
25. Select the ports that you want to provide connectivity to an IP multicast router.
26. Click **Ok**. The ports that are enabled appear in the SnoopMRouterPorts field.
27. Select **DynamicDowngradeEnable** check box to enable Ethernet Routing Switch 8800/8600 downgrades.
28. Select **CompatibilityModeEnable** check box to enable or disable v2-v3 compatibility mode.
29. Select the **SnoopQuerierEnable** check box to enable snoop querier.
30. Select the **SnoopQuerierAddr** check box and then type the IP address of the IGMP snoop querier.

31. Select the **MvrEnable** check box to enable or disable MCast-VLAN-registration on a VLAN.
32. Select the **MvrProxyEnable** check box to enable or disable MCast-VLAN-registration proxy on a VLAN.
33. Click **Apply**.

Variable definitions

Use the data in the following table to configure the IGMP tab.

Variable	Value
QueryInterval	The frequency (in seconds) at which the interface transmits IGMP host query packets. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	The maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds.) **Important:** You must configure this value lower than the QueryInterval.
Robustness	Use this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If a network is expected to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that one query for each query interval is dropped without the querier aging out.
LastMembQueryIntvl	The maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Avaya recommends that you configure this parameter to values greater than 3. If a fast leave process is not required, Avaya recommends that you use values greater than 10. (The value 3 is equal to 0.3 seconds, and 10 is equal to 1 second.)

Variable	Value
SnoopEnable	Enables or disables snoop. The default is disabled.
SsmSnoopEnable	Enables or disables support for PIM Source Specific multicast (SSM) on the snoop interface. The default is disabled.
ProxySnoopEnable	Enables or disables proxy snoop. The default is disabled.
IgapEnable	Enables or disables the Internet Group membership Authentication Protocol (IGAP) on this interface. The default is disabled.
AccntEnable	Enables or disables IGAP accounting on this interface. The default is enabled.
AuthEnable	Enables or disables IGAP authentication on this interface. The default is enabled.
Version	The version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables or disables fast leave on the interface. The default is disabled.
StreamLimitEnable	Enables or disables stream limitation on this VLAN.
Maximum Number of Streams	Specifies the maximum number of streams
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.
FastLeavePortMembers	Identifies the set of ports that are enabled for fast leave.
SnoopMRouterPorts	The set of ports in this interface that provide connectivity to an IP multicast router.
DynamicDowngradeEnable	Configures if the Avaya Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected, which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear, which means IGMPv3 is not compatible with IGMPv2.
SnoopQuerierEnable	Enables snoop querier.

Variable	Value
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
MvrEnable	Enables or disables MCast-VLAN-registration on a VLAN. It also specifies a VLAN works as a MCast-VLAN. A maximum of 16 MCast-VLAN are supported.
MvrProxyEnable	Enables or disables MCast-VLAN-registration proxy on a VLAN. You must enable MCast-VLAN-registration first.

Chapter 6: IP multicast basic configuration using the CLI

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts subscribe to multicast services using a host membership protocol. The Internet Group Management Protocol (IGMP) is an example of a host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast-Sparse Mode (PIM–SM) and Protocol Independent Multicast-Source Specific Multicast (PIM–SSM).

IP multicast basic configuration procedures

This task flow shows you the sequence of procedures you perform to configure basic elements of IP multicast routing.

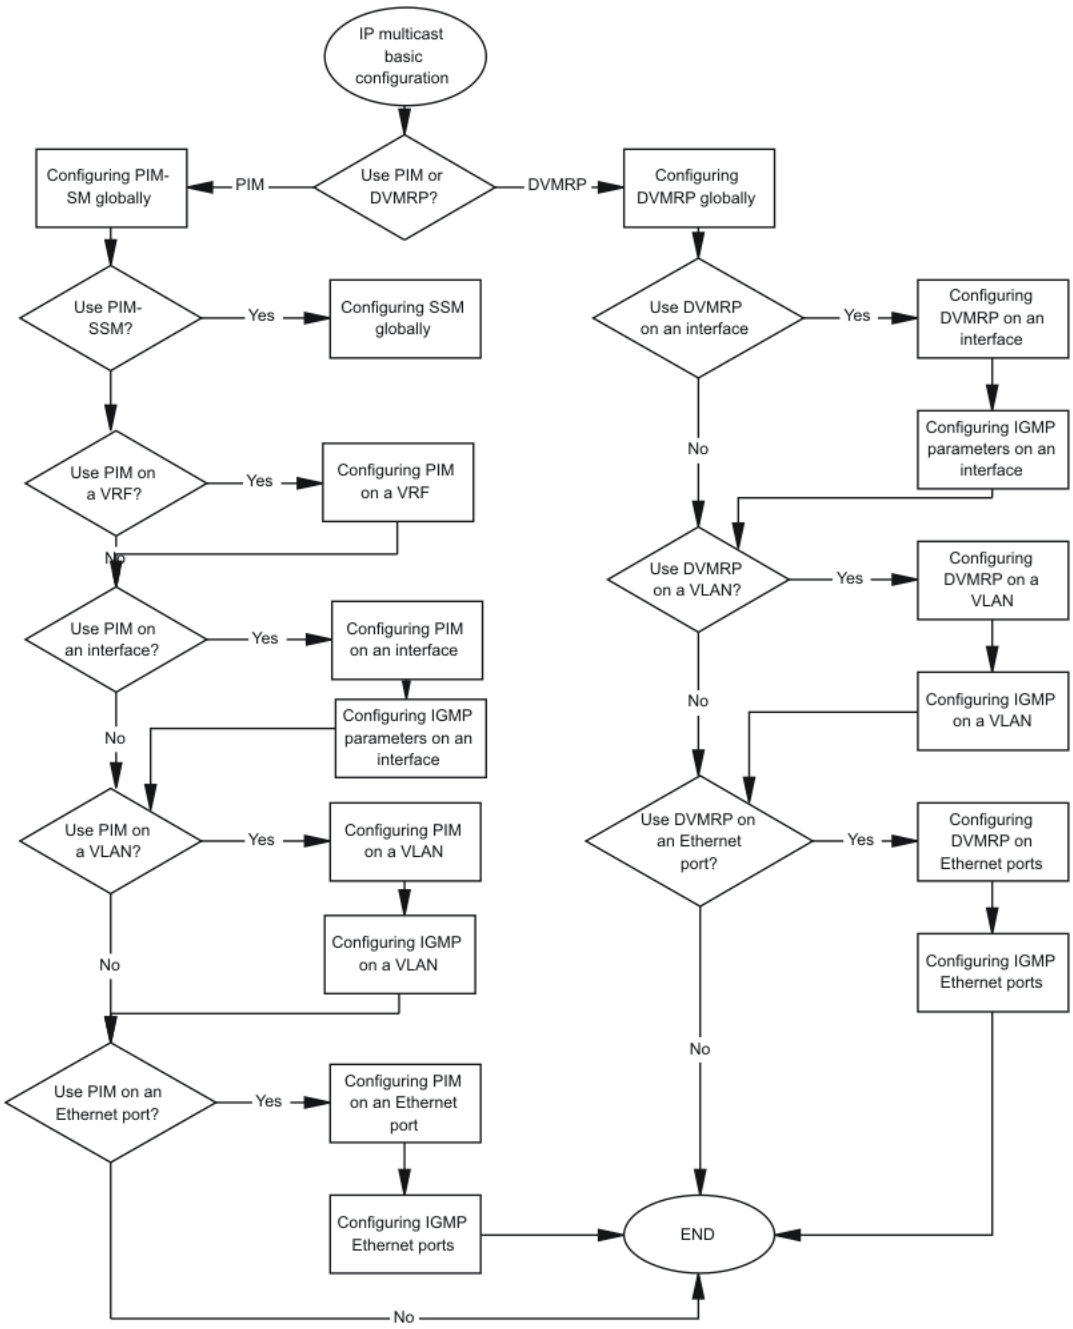

Figure 20: IP multicast basic configuration procedures

IP multicast basic configuration navigation

- Job aid on page 109
- Configuring PIM-SM globally on page 116
- Configuring PIM on a VRF on page 119
- Configuring PIM on an interface on page 121
- Configuring PIM on a VLAN on page 123
- Configuring PIM on an Ethernet port on page 125
- Configuring SSM globally on page 126
- Configuring DVMRP globally on page 127
- Configuring DVMRP on an interface on page 129
- Configuring DVMRP on a VLAN on page 130
- Configuring DVMRP on Ethernet ports on page 132
- Configuring IGMP on an interface on page 133
- Configuring IGMP on a VLAN on page 137
- Configuring IGMP Ethernet ports on page 141

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ethernet <ports> ip dvmrp`	`advertise-self <enable|disable>`
	`create <active|passive>`
	`default-listen <enable|disable>`
	`default-supply <enable|disable>`
	`default-supply-metric <cost>`

Command	Parameter
	disable
	enable
	in-policy <policy_name>
	info
	interface-type <active| passive>
	metric <cost>
	out-policy <policy_name>
config ethernet <ports> ip igmp	compatibility-mode enable <true|false>
	dynamic-downgrade-version enable <true|false>
	fast-leave <enable|disable>
	flush <mrouter|sender|grp-member> [<SenderAddress>] [<GroupAddress>]
	info
	last-memb-query-int <1/10_seconds>
	query-interval <seconds>
	query-max-resp <1/10_seconds>
	robustval <integer>
	router-alert <enable|disable>
	version <integer>
config ethernet <ports> ip pim	create <active|passive>
	disable
	enable
	hellointerval <seconds>
	info
	interface-type <active| passive>
	joinprune-interval <seconds>

Command	Parameter
config ip dvmrp	disable
	enable
	fwd-cache-timeout <integer>
	generate-log <enable|disable>
	generate-trap <enable| disable>
	info
	leaf-timeout <integer>
	nbr-probe-interval <integer>
	nbr-timeout <integer>
	route-discard-timeout <integer>
	route-expiration-timeout <integer>
	route-switch-timeout <integer>
	show-next-hop-table <enable| disable>
	triggered-update- interval <integer>
	update-interval <integer>
	prune-resend <enable|disable>
config ip dvmrp interface <ipaddr>	advertise-self <enable| disable>
	create <active|passive>
	default-listen <enable| disable>
	default-supply <enable| disable>
	default-supply-metric <cost>
	disable
	enable
	in-policy <policy_name>

Command	Parameter
	info
	interface-type <active| passive>
	metric <cost>
	out-policy <policy_name>
config ip igmp interface <ipaddr>	compatibility-mode enable <true|false>
	dynamic-downgrade-version enable <true|false>
	fast-leave <enable|disable>
	flush <mrouter|sender|grp-member> [<SenderAddress>] [<GroupAddress>]
	info
	last-memb-query-int <1/10_seconds>
	proxy-snoop <enable|disable>
	query-interval <seconds>
	query-max-resp <1/10_seconds>
	robustval <integer>
	router-alert <enable|disable>
	snoop <enable|disable>
	ssm-snoop <enable|disable>
	version <integer>
config ip pim	bootstrap-period <integer>
	c-rp-adv-timeout <integer>
	disable
	disc-data-timeout <integer>
	enable
	fwd-cache-timeout <integer>
	info
	joinprune-interval <integer>

Command	Parameter
	mode <sparse|ssm>
	register-suppression-timeout <integer>
	unicast-route-change-timeout <integer>
config ip pim interface <ipaddr>	create <active|passive>
	disable
	enable
	hellointerval <seconds>
	info
	interface-type <active| passive>
	joinprune-interval <seconds>
config ip vrf <vrfName> igmp interface <ipaddr>	fast-leave <enable|disable>
	flush <mrouter|sender|grp-member> [<SenderAddress>] [<GroupAddress>]
	info
	last-memb-query-int <1/10_seconds>
	multicast-vlan-registration <enable|disable>
	multicast-vlan-registration-proxy <enable|disable>
	proxy-snoop <enable|disable>
	query-interval <seconds>
	query-max-resp <1/10_seconds>
	robustval <integer>
	router-alert <enable|disable>
	snoop <enable|disable>
	snoop-querier-addr <ipaddr>
	snoop-querier <enable| disable>

Command	Parameter
	ssm-snoop <enable|disable>
	version <integer>
config ip vrf <vrfName> pim	bootstrap-period <integer>
	c-rp-adv-timeout <integer>
	create
	delete
	disable
	disc-data-timeout <integer>
	enable
	fast-joinprune <enable| disable>
	fwd-cache-timeout <integer>
	info
	joinprune-interval <integer>
	mode <sparse|ssm>
	register-suppression-timeout <integer>
	unicast-route-change-timeout <integer>
config ip vrf <vrfName> pim interface <ipaddr>	create <active|passive>
	disable
	enable
	hellointerval <seconds>
	info
	interface-type <active| passive>
	joinprune-interval <seconds>
config vlan <vid> ip dvmrp	advertise-self <enable| disable>
	create <active|passive>
	default-listen <enable| disable>

Command	Parameter
	default-supply <enable| disable>
	default-supply-metric <cost>
	disable
	enable
	in-policy <policy_name>
	info
	interface-type <active| passive>
	metric <cost>
	out-policy <policy_name>
config vlan <vid> ip igmp	del-mrouter <ports>
	fast-leave <enable|disable>
	flush <mrouter|sender| grp-member>[<SenderAddress>] [<GroupAddress>]
	info
	last-memb-query-int <1/10_seconds>
	mrouter <ports>
	multicast-vlan-registration <enable|disable>
	multicast-vlan-registration-proxy <enable|disable>
	proxy-snoop <enable|disable>
	query-interval <seconds>
	query-max-resp <1/10_seconds>
	robustval <integer>
	router-alert <enable|disable>
	snoop <enable|disable>
	snoop-querier-addr <ipaddr>

Command	Parameter
	snoop-querier <enable|disable>
	ssm-snoop <enable|disable>
	version <integer>
config vlan <vid> ip pim	**create <active|passive>**
	disable
	enable
	hellointerval <seconds>
	info
	interface-type <active|passive>
	joinprune-interval <seconds>
show ip pim info	
show vlan info pim	

Configuring PIM-SM globally

Configure PIM-SM to enable or disable PIM-SM globally on the switch and change default global parameters.

Procedure steps

1. Enable PIM-SM:

   ```
   config ip pim enable
   ```

2. Configure the remaining parameters as required.
3. Verify your configuration by displaying the global status of PIM on the switch:

   ```
   show ip pim info
   ```

Variable definitions

Use the data in the following table to use the **config ip pim** command.

Variable	Value
bootstrap-period <*integer*>	Specifies the interval (in seconds) that the elected bootstrap router (BSR) waits between originating bootstrap messages. • *integer* is an integer in the range of 5–32757. The default is 60.
c-rp-adv-timeout <*integer*>	Specifies how often (in seconds) a router configured as a candidate rendezvous point (C-RP) router sends C-RP advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected BSR. • *integer* is an integer in the range of 5–26214. The default is 60.
disable	Globally disables PIM on the switch.
disc-data-timeout <*integer*>	Specifies how long (in seconds) to discard data until the join is received from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the the timer expires or a join is received. • *integer* is an integer in the range of 5–65535. The default is 60.
enable	Globally activates PIM on the switch.
fwd-cache-timeout <*integer*>	Specifies the forward cache timeout value. • *integer* is an integer in the range of 10–86400. The default is 210. **Important:** When you configure one of the timers, activity-chk-interval or fwd-cache-timeout, with a nondefault value, you cannot configure the other timer.
info	Displays current PIM settings on the switch.
joinprune-interval <*integer*>	Specifies how long to wait (in seconds) before the PIM router sends out the next join/prune message to its upstream neighbors. • *integer* is an integer in the range of 1–18724. The default is 60.
mode <*sparse*|*ssm*>	Configures the mode of this interface globally. After you change from one mode to another, an information message appears to remind you that traffic does not stop immediately. The default is sparse.
register-suppression-timeout <*integer*>	Specifies how long (in seconds) the designated router (DR) suppresses sending registers to the RP. The timer starts

Variable	Value
	after the DR receives a register-stop message from the RP. • *integer* is an integer in the range of 6–65535. The default is 60.
unicast-route-change-timeout *<integer>*	Specifies how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates for PIM. **Important:** Lowering this value increases how often the switch polls the RTM. This can affect the performance of the switch, especially when a high volume of traffic flows through the switch. • *integer* is an integer in the range of 2–65535. The default is 5.

Job aid

The following table shows the field descriptions for this command.

Field	Description
PimStat	Indicates the status of PIM.
Mode	Indicates the PIM mode.
Mbr	Indicates the status of the PIM multicast border router feature.
StaticRP	Indicates the status of static RP.
BootstrapPeriod	Indicates the interval between originating bootstrap messages at the elected BSR.
CRPAdvTimeout	Indicates the candidate RP timer (in seconds) for sending C-RP-Adv messages.
DiscDataTimeout	Indicates the time (in seconds) used to discard data until the join is received from the RP. An IP multicast discard record is created and deleted after the timer expires and after a join is received.
FwdCacheTimeout	Indicates the PIM forward cache expiry value in seconds. This value is used in aging PIM mroutes.
RegSupprTimeout	Indicates the Register-Suppression timer in seconds.
UniRouteChangeTimeout	Indicates the frequency at which the RTM is polled for routing information updates.

Field	Description
JoinPruneInt	Indicates the join pruning interval in seconds.

Configuring PIM on a VRF

Configure PIM to create or remove an instance of the multicast routing protocol on a Virtual Router Forwarding (VRF) instance. Use the VRF Lite feature with multicast routing protocols to create multiple virtual multicast routers. By default, PIM does not run on a VRF.

Prerequisites

- Multicast virtualization works on an R, RS, or 8800 module port.
- Multicast virtualization requires an 8692 SF/CPU with SuperMezz or an 8895 SF/CPU.
- You must configure virtualization for the unicast routing system.

Procedure steps

1. Create a PIM instance on a VRF:

   ```
   config ip vrf <vrfName> pim create
   ```

2. Remove a PIM instance from a VRF:

   ```
   config ip vrf <vrfName> pim delete
   ```

Variable definitions

Use the data in the following table to use the **`config ip vrf pim`** command.

Variable	Value
bootstrap-period <integer>	Specifies the interval (in seconds) that the elected bootstrap router (BSR) waits between originating bootstrap messages. `integer` is an integer in the range of 5–32757. The default is 60.
c-rp-adv-timeout <integer>	Specifies how often (in seconds) a router configured as a candidate rendezvous point

Variable	Value
	(C-RP) router sends C-RP advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected BSR. `integer` is an integer in the range of 5–26214. The default is 60.
create	Creates a PIM instance on the VRF.
delete	Removes a PIM instance from the VRF.
disable	Globally disables PIM on the VRF.
disc-data-timeout <integer>	Specifies how long (in seconds) to discard data until the join is received from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the the timer expires or a join is received. `integer` is an integer in the range of 5–65535. The default is 60.
enable	Activates the PIM instances on the VRF.
fast-joinprune <enable|disable>	Activates or deactivates the fast join prune feature. The default is disable.
fwd-cache-timeout <integer>	Specifies the forward cache timeout value. `integer` is an integer in the range of 10–86400. The default is 210.
info	Displays current PIM settings on the VRF.
joinprune-interval <integer>	Specifies how long to wait (in seconds) before the PIM router sends out the next join/prune message to its upstream neighbors. `integer` is an integer in the range of 1–18724. The default is 60.
mode <sparse|ssm>	Configures the mode of this interface globally. After you change from one mode to another, an information message appears to remind you that traffic does not stop immediately. The default is sparse.
register-suppression-timeout <integer>	Specifies how long (in seconds) the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP.

Variable	Value
	`integer` is an integer in the range of 6–65535. The default is 60.
unicast-route-change-timeout <integer>	Specifies how often (in seconds) the VRF polls the routing table manager (RTM) for unicast routing information updates for PIM. `integer` is an integer in the range of 2–65535. The default is 5. **Important:** Lowering this value increases how often the VRF polls the RTM. This can affect the performance of the switch, especially when a high volume of traffic flows through the switch.
vrfName	Specifies the name of the VRF instance, a string length from 0–16 characters.

Configuring PIM on an interface

Configure PIM for each interface to enable the interface to perform multicasting operations.

You configure PIM on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Prerequisites

- You must enable PIM globally before you configure it on an interface.

Procedure steps

1. Enable PIM on an interface:

   ```
   config ip pim interface <ipaddr> enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **config ip pim interface** and **config ip vrf <vrfName> pim interface** commands.

Variable	Value
create <active|passive>	Enables PIM on a specific interface with a specific type. An active interface accepts PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful when a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.
disable	Disables PIM on the local switch interface.
enable	Enables PIM on the local switch interface.
hellointerval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default is 30 seconds.
info	Displays current PIM configuration settings on the local switch interface.
interface-type <active| passive>	Specifies whether the selected interface is active or passive. You can change the state of a PIM interface after you create the interface but only if you disable PIM on the interface. The default is active.
ipaddr	Indicates the IP address of the selected interface.
joinprune-interval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default is 60 seconds.

Example of configuring PIM on an interface

Procedure steps

1. Create a PIM passive interface:

   ERS-8606:5# **config ip pim interface 128.3.1.1 create passive**

2. Display information about the interface:

   ERS-8606:5/config/ip/pim/interface/128.3.1.1# **info**

```
Sub-Context:
Current Context:
```

```
enable : true
mode : sparse
interface-type : passive
joinpruneint : 60
cbsrpref : -1 (disabled)
```

Configuring PIM on a VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable PIM globally before you configure PIM on a VLAN.

Procedure steps

1. Enable PIM on a VLAN:

   ```
   config vlan <vid> ip pim enable
   ```

2. Configure the remaining parameters as required.
3. Verify your configuration by displaying information about the PIM-SM interface setup for VLANs:

   ```
   show vlan info pim
   ```

Variable definitions

Use the data in the following table to use the **`config vlan ip pim`** command.

Variable	Value
create <active|passive>	Enables PIM on a specific VLAN with a specific type. • **`active`** accepts PIM control transmitted and received traffic. • **`passive`** prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful when a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.

Variable	Value
disable	Disables PIM on the selected VLAN.
enable	Enables PIM on the selected VLAN.
hellointerval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default is 30 seconds.
info	Displays current PIM configuration settings on the selected VLAN.
interface-type <active| passive>	Specifies whether the selected interface is active or passive. You can change the state of a PIM interface after you create it but only if you disable PIM on the interface.
joinprune-interval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default is 60 seconds.
vid	Identifies a VLAN ID from 1–4092.

Job aid

The following table shows the field descriptions for this command.

Field	Description
VLAN-ID	Identifies the VLAN.
PIM-ENABLE	The state of PIM on the VLAN.
MODE	The configured mode of this VLAN. The valid modes are SSM and Sparse.
HELLOINT	Indicates how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Indicates how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default join/prune interval is 60 seconds.
CBSR PREF	The preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
INTF TYPE	Indicates whether the PIM interface is active or passive.

Configuring PIM on an Ethernet port

Configure PIM for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable PIM globally before you configure PIM on an Ethernet port.

Procedure steps

1. Enable PIM on an Ethernet port:

   ```
   config ethernet <ports> ip pim enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ethernet ip pim`** command.

Variable	Value
create <active|passive>	Enables PIM on a specific brouter port with a specific type. An active port accepts PIM control transmitted or received traffic. A passive port prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful when a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.
disable	Disables PIM on the selected brouter port.
enable	Enables PIM on the selected port.
hellointerval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default is 30 seconds.
info	Displays current PIM configuration settings on the selected brouter port.
interface-type <active| passive>	Specifies whether the selected port is active or passive. You can change the state of a PIM interface after you

Variable	Value
	create it but only if you disable PIM on the port. The default is active.
joinprune-interval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default is 60 seconds.
ports	Specifies the port using the convention {slot/port[-slot/port][,...]}.

Configuring SSM globally

Configure SSM to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces running PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Prerequisites

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shorted Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP and OSPF, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — OSPF and RIP,* (NN46205-522).

  PIM requires a unicast protocol to forward multicast traffic within the network when it performs the Reverse Path Forwarding (RFP) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that activates PIM-enabled routers to communicate. The unicast routing table must contain a route to every multicast source in the network as well as routes to PIM entities like the RPs and BSR.
- Enable PIM globally.

Procedure steps

Configure PIM-SSM:

```
config ip pim mode ssm
```

The following message appears:

```
WARNING: All Static Source Group entries in the SSM range will
be deleted

Do you wish to continue? (y/n) ?
```

Configuring DVMRP globally

Configure the Distance Vector Multicast Routing Protocol (DVMRP) globally to enable or disable DVMRP and change default global parameters.

Procedure steps

1. Enable DVMRP globally:

   ```
   config ip dvmrp enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ip dvmrp`** command.

Variable	Value
disable	Globally disables DVMRP on the switch.
enable	Globally activates DVMRP on the switch.
fwd-cache-timeout <*integer*>	Configures the forward cache timeout (in seconds). • *integer* is the range of 300–86400 seconds. The default value is 300 seconds.
generate-log <enable|disable>	Enables or disables the DVMRP log. The default is disable.
generate-trap <enable|disable>	Enables or disables the DVMRP trap. The default is disable.
info	Displays DVMRP settings on the switch.

Variable	Value
leaf-timeout <integer>	Configures the length of time (in seconds) the router waits for a response from a neighbor before considering the attached network as a leaf network. • *integer* is the range of 25–4000 seconds. The default value is 125 seconds.
nbr-probe-interval <integer>	Configures the time interval (in seconds) for the DVMRP router to send a neighbor probe message on its interface. • *integer* is the range of 5–30 seconds. The default value is 10 seconds.
nbr-timeout <*integer*>	Configures the length of time (in seconds) the router waits to receive a report from a neighbor before considering the connection inactive. • *integer* is the range of 35–8000 seconds. The default value is 35 seconds.
prune-resend <enable|disable>	Sends prune messages every 3 minutes, to address the link failures at remote upstream switches. The feature is disabled by default.
route-discard-timeout <*integer*>	Configures the route discard timeout (in seconds). • *integer* is the range of 40–8000. The default value is 260 seconds.
route-expiration-timeout <*integer*>	Configures the route expiration timeout (in seconds). • *integer* is the range of 20–400 seconds. The default value is 140 seconds.
route-switch-timeout <*integer*>	Configures the route switch timeout (in seconds). • *integer* is the range of 20–2000. The default value is 140 seconds.
show-next-hop-table <enable|disable>	Enables or disables showing information about the DVMRP next hops. The default is disable.
triggered-update-interval <*integer*>	Configures the time interval (in seconds) between triggered update messages sent after routing information changes. • *integer* is the range of 5–1000 seconds. The default value is 5 seconds.
update-interval <*integer*>	Configures the time interval (in seconds) between DVMRP router update messages. • *integer* is the range of 10–2000 seconds. The default value is 60 seconds.

Configuring DVMRP on an interface

Configure DVMRP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on an interface.

Procedure steps

1. Enable DVMRP on an interface:

   ```
   config ip dvmrp interface <ipaddr> enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ip dvmrp interface`** command.

Variable	Value
advertise-self <enable|disable>	Enables or disables the advertisement of local routes for the selected interface to other switches in the network. The default is enable.
create <active|passive>	Configures the interface type: active or passive. The default is active.
default-listen <enable|disable>	Learns the default route over the specified interface if this feature is enabled on the interface. The default setting is enable.
default-supply <enable|disable>	Generates and advertises the default route if this feature is enabled on the interface. The default setting is disable.
default-supply-metric <*cost*>	Advertises the specified metric over the interface if you configured the interface to supply the default route. The range is 1–31 hops. The default setting is 1 hop.

Variable	Value
disable	Disables DVMRP on the local router interface.
enable	Enables DVMRP on the local router interface.
in-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64 characters.
info	Displays information about the specified DVMRP local router interface.
interface-type <active|passive>	Configures an interface as active or passive. The default is active.
ipaddr	Indicates the IP address of the selected interface.
metric <*cost*>	Configures the cost metric (maximum number of hops) for the router interface. • *cost* is the range of 1–31. The default is 1.
out-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64 characters.

Configuring DVMRP on a VLAN

Configure DVMRP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on a VLAN.

Procedure steps

1. Enable DVMRP on a VLAN:

   ```
   config vlan <vid> ip dvmrp enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip dvmrp`** command.

Variable	Value
advertise-self <enable|disable>	Enables or disables the advertisement of local routes for the selected VLAN to other switches in the network. The default is enable.
create <active|passive>	Configures the interface type: active or passive. The default is active.
default-listen <enable|disable>	Learns the default route over the specified VLAN if this feature is enabled on the interface. The default setting is enable.
default-supply <enable|disable>	Generates and advertises the default route if this feature is enabled on the VLAN. The default setting is disable.
default-supply-metric <*cost*>	Advertises the specified metric over the VLAN if you configure the VLAN to supply the default route. The range is 1–31 hops. The default setting is 1 hop.
disable	Disables DVMRP on the VLAN.
enable	Enables DVMRP on the VLAN.
in-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64 characters.
info	Displays DVMRP settings on the VLAN.
interface-type <active|passive>	Configures an interface as active or passive. The default is active.
metric <*cost*>	Configures the DVMRP route metric. • *cost* is the maximum number of hops with a value of 1–31. The default is 1.
out-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64.
vid	Specifies a VLAN ID from 1–4092.

Configuring DVMRP on Ethernet ports

Configure DVMRP on each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on an Ethernet port.

Procedure steps

1. Enable DVMRP on an Ethernet port:

   ```
   config ethernet <ports> ip dvmrp enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ethernet ip dvmrp`** command.

Variable	Value
advertise-self <enable|disable>	Enables or disables the advertisement of local routes for the selected port to other switches in the network. The default is enable.
create <active|passive>	Configures the interface type: active or passive. The default is active.
default-listen <enable|disable>	Learns the default route over the specified port if this feature is enabled on the interface. The default setting is enable.
default-supply <enable|disable>	Generates and advertises the default route if this feature is enabled on the port. The default setting is disable.
default-supply-metric <*cost*>	Advertises the specified metric over the port if you use configured the port to supply the default route. The range is 1–31 hops. The default setting is 1 hop.

Variable	Value
disable	Disables DVMRP on the port.
enable	Enables DVMRP on the port.
in-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64 characters.
info	Displays DVMRP settings on the port.
interface-type <active|passive>	Configures an interface as active or passive. The default is active.
metric <*cost*>	Configures the DVMRP route metric. • *cost* is the maximum number of hops with a value of 1–31. The default is 1.
out-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64 characters.
ports	Identifies the port using the convention {slot/port[-slot/port][,...]}.

Configuring IGMP on an interface

Configure IGMP for each interface to change default multicasting operations. You can specify the version of IGMP and backward compatibility operations.

You configure IGMP on a VRF the same way you configure IGMP for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Prerequisites

- You must globally enable either PIM or DVMRP.
- You must enable either PIM or DVMRP on the interface.

Procedure steps

1. Configure IGMP on an interface:

   ```
   config ip igmp interface <ipaddr>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ip igmp interface`** and **`config ip vrf <vrfName> igmp interface`** commands.

Variable	Value
compatibility-mode enable <true|false>	Enables or disables v2-v3 compatibility mode. The default value is false, which means IGMPv3 is not compatible with IGMPv2.
dynamic-downgrade-version enable <true|false>	Configures if the Avaya Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is true, which means the switch downgrades to the oldest version of IGMP on the network.
fast-leave <enable|disable>	Enables or disables the fast leave option on the interface. The default is disable.
flush <mrouter|sender|grp-member> [<SenderAddress>] [<GroupAddress>]	Flushes the specified table.
info	Displays the access list of the IGMP interface.
ipaddr	Indicates the IP address of the selected interface.
last-memb-query-int *<1/10_seconds>*	The maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. • *1/10_seconds* is an integer in the range from 0–255, and the default is 10 tenths of a second. Avaya

Variable	Value
	recommends that you configure this value between 3–10 (equal to 0.3 – 1.0 seconds).
proxy-snoop <enable|disable>	Enables or disables the Layer 3 proxy-snoop option. The default is disable.
query-interval <*seconds*>	Configures the frequency (in seconds) at which the interface transmits host query packets. • *seconds* is an integer in the range from 1–65535 with a default of 125.
query-max-resp <*1/10_seconds*>	The maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. • *1/10_seconds* is an integer value with a range of 0–255, and the default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robustval <*integer*>	Use this variable to tune for the expected packet loss of a network. • *integer* is an integer value with a range of 2–255 seconds. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.
router-alert <enable|disable>	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set. **Important:** To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
snoop <enable|disable>	Enables or disables the Layer 3 snoop option. The default is disable.

Variable	Value
ssm-snoop <enable|disable>	Enables or disables support for PIM-SSM on the snoop interface. **Important:** SSM-snoop assumes that only one source for each group exists in the network. Traffic for all sources for a group are permitted on a port that receives an IGMPv3 join. The default is disable.
version <*integer*>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. • *integer* is an integer value with a value of 1, 2, or 3. The default value is 2 (IGMPv2).

Example of configuring IGMP on an interface

Procedure steps

1. Set the last member query interval to 15 tenths of a second (equal to 1.5 seconds).

```
ERS-8606:5#config ip igmp interface 10.10.99.151 last-memb-query-int 15
```

2. Set the query interval to 100 seconds.

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151# query-interval 100
```

3. Set the query maximum response time to 15 tenths of a second (equal to 1.5 seconds).

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151# query-max-resp 50
```

4. Set the robustness value to 4 seconds.

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151#robustval 4
```

5. Enable IGMPv3.

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151# version 3
```

6. Enable the fast leave option.

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151#fast-leave enable
```

7. Enable support for SSM on the snoop interface.

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151# ssm-snoop enable
```

8. Use the **`info`** command to view your configuration.

```
ERS-8606:5/config/ip/igmp/interface/10.10.99.151# info
```

```
Sub-Context: access-control igap mrdisc static-members stream
-limit-members dynamic-downgrade-version compatibility-mode
Current Context:
last-memb-query-int : 15
query-interval : 100 secs
query-max-resp : 50
robustval : 4
version : 3
fast-leave : enable
ssm-snoop  : enable
```

Configuring IGMP on a VLAN

Configure IGMP for each interface to change default multicasting operations. You can specify the version of IGMP and backward compatibility operations.

Prerequisites

- You must globally enable either PIM or DVMRP.
- You must enable either PIM or DVMRP on the VLAN.

Procedure steps

Configure IGMP on a VLAN:

```
config vlan <vid> ip igmp
```

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp`** command.

Variable	Value
del-mrouter *<ports>*	Deletes multicast router ports.
fast-leave <enable|disable>	Prevents a port from receiving a leave message from a member of a group. Normal IGMP behavior is skipped. The default is disable.
flush <mrouter|sender|grp-member> *[<SenderAddress>] [<GroupAddress>]*	Flushes the specified table.
info	Displays IGMP settings on the VLAN.

Variable	Value
last-memb-query-int *<1/10_seconds>*	The maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. • *1/10_seconds* is the range from 0–255, and the default is 10 tenths of a second. Avaya recommends that you configure this value between 3 and 10 (equal to 0.3 – 1.0 seconds).
mrouter *<ports>*	Adds multicast router ports.
multicast-vlan-registration <enable|disable>	Enables or disables the multicast-vlan-registration option globally for the VLAN. The default is disable.
multicast-vlan-registration-proxy <enable|disable>	Enables or disables the multicast-vlan-registration-proxy option globally for the VLAN. The default is disable.
proxy-snoop <enable|disable>	Enables or disables the proxy-snoop option globally for the VLAN. The default is disable.
query-interval *<seconds>*	Configures the frequency (in seconds) at which the host query packets are transmitted on the VLAN. • *seconds* is the range from 1–65535. The default value is 125 seconds.
query-max-resp *<1/10_seconds>*	The maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. • *1/10_seconds* is an integer value with a range of 0–255, and the default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robustval <integer>	Use this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If a network is expected to lose query packets, increase the robustness value. • *integer* is an integer value with a range of 2–255. The default value is 2. The default value of 2 means that

Variable	Value
	one query for each query interval is dropped without the querier aging out.
router-alert <enable|disable>	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set. **Important:** To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
snoop <enable|disable>	Enables or disables the snoop option for the VLAN. The default is disable.
igmp snoop-querier-addr <ipaddr>	Enables or disables support for PIM Source Specific multicast (SSM) on the snoop interface. The default is disable.
snoop-querier <enable|disable>	Enables or disables snoop querier. The default is disable.
ssm-snoop <enable|disable>	Enables or disables support for PIM-SSM on the snoop interface. **Important:** SSm-snoop assumes that only one source for each group exist in the network. Traffic for all sources for a group are permitted on a port that receives an IGMP v3 join.
version <integer>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. • *integer* is an integer value with a value of 1, 2, or 3. The default value is 2 (IGMPv2).

Example of configuring IGMP on a VLAN

Procedure steps

1. Set the last member query interval to 15 tenths of a second (equal to 1.5 seconds).

```
ERS-8606:5#config vlan 2 ip igmp last-memb-query-int 15
```

2. Set the query interval to 100 seconds.

```
ERS-8606:5/config/vlan/2/ip/igmp#query-interval 100
```

3. Set the query maximum response time to 15 tenths of a second (equal to 1.5 seconds).

```
ERS-8606:5/config/vlan/2/ip/igmp# query-max-resp 50
```

4. Set the robustness variable to 4.

```
ERS-8606:5/config/vlan/2/ip/igmp# robustval 4
```

5. Enable IGMPv3.

```
ERS-8606:5/config/vlan/2/ip/igmp# version 3
```

6. Enable proxy snoop for the VLAN.

```
ERS-8606:5/config/vlan/2/ip/igmp# proxy-snoop enable
```

7. Enable snoop for the VLAN.

```
ERS-8606:5/config/vlan/2/ip/igmp# snoop enable
```

8. Enable support for SSM on the snoop interface.

```
ERS-8606:5/config/vlan/2/ip/igmp# ssm-snoop enable
```

9. Enable the fast leave option.

```
ERS-8606:5/config/vlan/2/ip/igmp# fast-leave enable
```

10. Use the **`info`** command to view your configuration.

```
8610co:5/config/vlan/2/ip/igmp# info
```

```
Sub-Context: access-control fast-leave-members igap mrdisc static-members
stream
-limit-members dynamic-downgrade-version compatibility-mode

Current Context:
last-memb-query-int : 15
query-interval : 100 secs
query-max-resp : 50
robustval : 4
version : 3
proxy-snoop  : enable
snoop  : enable
mrouter :
ssm-snoop  : enable
fast-leave : enable
router-alert : enable
```

Configuring IGMP Ethernet ports

Configure IGMP for each interface to change default multicasting operations. You can specify the version of IGMP and backward compatibility operations.

Prerequisites

- You must globally enable either PIM or DVMRP.
- You must enable either PIM or DVMRP on the port.

Procedure steps

Configure IGMP on Ethernet ports:

```
config ethernet <ports> ip igmp
```

Variable definitions

Use the data in the following table to use the **`config ethernet ip igmp`** command.

Variable	Value
compatibility-mode enable <true|false>	Enables or disables v2-v3 compatibility mode. The default value is false, which means IGMPv3 is not compatible with IGMPv2.
dynamic-downgrade-version enable <true|false>	Configures if the Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is true, which means the switch downgrades to the oldest version of IGMP on the network.
fast-leave <enable|disable>	Enables or disables the fast leave mode, which a switch uses to immediately stop forwarding traffic for a multicast group as soon as an IGMPv2 leave group message is received on an interface. The default is disable.

Variable	Value
flush <mrouter|sender|grp-member> *[<SenderAddress>] [<GroupAddress>]*	Flushes the specified table.
info	Displays IGMP settings on the port.
last-memb-query-int *<1/10_seconds>*	The maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. • *1/10_seconds* is the range from 0–255, and the default is 10 tenths of a second. Avaya recommends that you configure this value between 3–10 (equal to 0.3 – 1.0 seconds).
query-interval *<seconds>*	Configures the frequency (in seconds) at which the port transmits host query packets. • *seconds* is the range of 1–65535 seconds. The default value is 125 seconds.
query-max-resp *<1/10_seconds>*	The maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. • *1/10_seconds* is an integer value with a range of 0–255 and a default of 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
ports	Specifies the port using the convention {slot/port[-slot/port][, ...]}.
robustval *<integer>*	Configures the expected packet loss of a network. • *integer* is the range of 2–255 with a default of 2. Increase the value if you expect the network to experience packet loss.
router-alert <enable|disable>	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set or not.

Variable	Value
	Important: To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
version <*integer*>	Configures the version of IGMP that you want to configure on this port. For IGMP to function correctly, all routers on a LAN must use the same version. • *integer* is an integer value with a value of 1, 2, or 3. The default value is 2 (IGMPv2).

Deleting a single IP Multicast record

Use this procedure to delete a single IP Multicast record.

Procedure steps

Delete a multicast record:

```
clear ip mroute group <ipaddr> vlan <vid> [source <ipaddr>]
```

Variable definitions

Use the data in the following table to use the **`clear ip mroute group`** command.

Variable	Value
group <ipaddr>	Specifies the multicast group address.
vlan <vid>	Specifies the VLAN ID.
[source <ipaddr>]	Specifies the multicast source address.

Chapter 7: IP multicast basic configuration using the ACLI

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts subscribe to multicast services using a host membership protocol. The Internet Group Management Protocol (IGMP) is an example of a host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast-Sparse Mode (PIM–SM) and Protocol Independent Multicast-Source Specific Multicast (PIM–SSM).

IP multicast basic configuration procedures

This task flow shows you the sequence of procedures you perform to configure basic elements of IP multicast routing. To link to a procedure, go to IP multicast basic configuration navigation on page 146.

Figure 21: IP multicast basic configuration procedures

IP multicast basic configuration navigation

- Job aid on page 147
- Configuring PIM-SM globally on page 151
- Configuring PIM on a VRF on page 155

- Configuring PIM on a VLAN on page 158
- Configuring PIM on an Ethernet port on page 160
- Configuring SSM globally on page 161
- Configuring DVMRP globally on page 162
- Configuring DVMRP on a VLAN on page 165
- Configuring DVMRP on Ethernet ports on page 167
- Configuring IGMP on a VLAN on page 169
- Configuring IGMP Ethernet ports on page 173

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	
show ip pim [vrf Word<0-16>] [vrfids Word<0-255>]	
show ip pim interface [vlan] [vrf Word<0-16>] [vrfids Word<0-255>]	
Global Configuration mode	
ip dvmrp	**fwd-cache-timeout <integer>**
	generate-log
	generate-trap
	leaf-timeout <integer>
	nbr-probe-interval <integer>
	nbr-timeout <integer>
	output-report-delay <integer>
	prune-resend
	route-discard-timeout <integer>
	route-expiration-timeout <integer>

Command	Parameter
	route-switch-timeout <integer>
	show-next-hop-table
	triggered-update-interval <integer>
ip pim	bootstrap-period <5-32757>
	disc-data-timeout <5-65535>
	enable
	fast-joinprune
	fwd-cache-timeout <10-86400>
	joinprune-interval <1-18724>
	mbr
	mode <sparse|ssm>
	register-suppression-timeout <6-65535>
	rp-c-adv-timeout <5-26214>
	rp-candidate group <A.B.C.D> <mask address> rp <A.B.C.D>
	static-rp <A.B.C.D/X> <A.B.C.D>
	static-rp specific-route
	unicast-route-change-timeout <2-65535>
	virtual-neighbor <A.B.C.D> <A.B.C.D>
Interface Configuration mode	
ip dvmrp	active
	advertise-self
	default-listen
	default-supply
	default-supply-metric <integer>
	enable
	in-policy <Word 0-64>

Command	Parameter
	interface-type <active>
	metric <integer>
	out-policy <Word 0-64>
ip igmp	**access-list**
	compatibility-mode
	dynamic-downgrade-version
	igap (VLAN Interface Configuration mode only)
	immediate-leave
	immediate-leave-members <portList> (VLAN Interface Configuration mode only)
	last-member-query-interval <integer>
	mrdisc (VLAN Interface Configuration mode only)
	mrouter <portList> (VLAN Interface Configuration mode only)
	port <portList> (GigabitEthernet or FastEthernet Interface Configuration Mode only)
	proxy (VLAN Interface Configuration mode only)
	query-interval <integer>
	query-max-response <integer>
	robust-value <integer>
	router-alert
	snooping (VLAN Interface Configuration mode only)
	ssm-snoop (VLAN Interface Configuration mode only)
	static-group (VLAN Interface Configuration mode only)
	stream-limit

Command	Parameter
	stream-limit-group (VLAN Interface Configuration mode only)
	stream-limit-max-streams
	version <integer>
ip pim	**active**
	bsr-candidate preference <integer>
	enable
	interface-type <active|passive>
	joinprune-interval <integer>
	passive
	query-interval <integer>
VRF Router Configuration mode	
ip pim	**bootstrap-period <5-32757>**
	disc-data-timeout <5–65535>
	enable
	fast-joinprune
	fwd-cache-timeout <10–8640>
	joinprune-interval <1–18724>
	mode <sparse|ssm>
	register-suppression-timeout <6–65535>
	rp-c-adv-timeout <5–26214>
	rp-candidate group <A.B.C.D> <mask address> rp <A.B.C.D>
	static-rp <A.B.C.D/X> rp <A.B.C.D>
	static-rp specific-route
	unicast-route-change-timeout <2–65535>
	virtual-neighbor <A.B.C.D> <A.B.C.D>

Configuring PIM-SM globally

Configure PIM-SM to enable or disable PIM-SM globally on the switch and change default global parameters.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable PIM-SM:

   ```
   ip pim enable
   ```

2. Configure the remaining parameters as required.
3. Verify your configuration by displaying the global status of PIM on the switch:

   ```
   show ip pim [vrf Word<0-16>] [vrfids Word<0-255>]
   ```

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
bootstrap-period <*5–32757*>	Specifies the interval (in seconds) that the elected bootstrap router (BSR) waits between originating bootstrap messages. • *integer* is an integer in the range of 5–32757. The default is 60. To set this option to the default value, use the **`default`** operator with the command.
disc-data-timeout <*5–65535*>	Specifies how long (in seconds) to discard data until the join is received from the rendezvous point (RP). An IP

Variable	Value
	multicast discard record is created after a register packet is sent, until the the timer expires or a join is received. • *integer* is an integer in the range of 5–65535. The default is 60. To set this option to the default value, use the **default** operator with the command.
enable	Globally activates PIM on the switch. To set this option to the default value, use the **default** operator with the command. The default is disabled.
fast-joinprune	Enables the fast join prune interval. To set this option to the default value, use the **default** operator with the command. The default is disabled.
fwd-cache-timeout <*10–86400*>	Specifies the forward cache timeout value. • *integer* is an integer in the range of 10–86400. The default is 210. **Important:** When you configure one of the timers, activity-chk-interval or fwd-cache-timeout, with a nondefault value, you cannot configure the other timer. To set this option to the default value, use the **default** operator with the command.
joinprune-interval <*1–18724*>	Specifies how long to wait (in seconds) before the PIM router sends out the next join/prune message to its upstream neighbors. • *integer* is an integer in the range of 1–18724. The default is 60. To set this option to the default value, use the **default** operator with the command.
mbr	Enables or disables the PIM multicast border router globally. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
mode <*sparse*|*ssm*>	Configures the mode of this interface globally. After you change from one mode to another, an information message appears to remind you that traffic does not stop immediately. To set this option to the default value, use the **default** operator with the command. The default is sparse.

Variable	Value
register-suppression-timeout *<6–65535>*	Specifies how long (in seconds) the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. • *integer* is an integer in the range of 6–65535. The default is 60. To set this option to the default value, use the **default** operator with the command.
rp-c-adv-timeout *<5–26214>*	Specifies how often (in seconds) a router configured as a candidate RP (C-RP) sends C-RP advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected BSR. • *integer* is an integer in the range of 5–26214. The default is 60. To set this option to the default value, use the **default** operator with the command.
rp-candidate group <A.B.C.D> <mask address> rp <A.B.C.D>	Adds or deletes candidate RP entries. Use the **no** operator to later remove this configuration.
static-rp <A.B.C.D/X> <A.B.C.D>	Adds static RP entries and activates static RP. A.B.C.D/X represents the group address and mask. A.B.C.D is the RP IP address.
static-rp specific-route	With static RP enabled, if the route to the RP is removed, the Avaya Ethernet Routing Switch 8800/8600 can fail over to an alternate static RP. However, if a default route exists in the routing table, that default route still appears as an active route to the failed RP. In this case, the switch does not fail over to the alternate RP. A similar situation exists with SMLT-based configurations, where an internal-only default static route is used during IST failover and recovery. In this case, the internal default route appears as an active route to the failed RP, and therefore does not failover to the alternate RP. To resolve these situations, you can configure the lookup for static RP to be chosen from the specific route rather than the best route. In this case, when the route to the active RP fails, the switch no longer interprets the default route as a valid route for RP purposes, and therefore fails over to the alternate RP.
unicast-route-change-timeout *<2–65535>*	Specifies how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates for PIM.

Variable	Value
	Important: Lowering this value increases how often the switch polls the RTM. This can affect the performance of the switch, especially when a high volume of traffic flows through the switch. • *integer* is an integer in the range of 2–65535. The default is 5. To set this option to the default value, use the **`default`** operator with the command.
virtual-neighbor <A.B.C.D> <A.B.C.D>	Adds a virtual neighbor to an interface globally. A.B.C.D represents the IP addresses of the interface and the virtual neighbor. Use the **`no`** operator to later remove this configuration.

Job aid

The following table shows the field descriptions for this command.

Field	Description
PimStat	Indicates the status of PIM.
Mode	Indicates the PIM mode.
Mbr	Indicates the status of the PIM multicast border router feature.
StaticRP	Indicates the status of static RP.
FastJoinPrune	Indicated the status of the fast join prune interval.
BootstrapPeriod	Indicates the interval between originating bootstrap messages at the elected BSR.
CRPAdvTimeout	Indicates the candidate RP timer (in seconds) for sending C-RP-Adv messages.
DiscDataTimeout	Indicates the time (in seconds) used to discard data until the join is received from the RP. An IP multicast discard record is created and deleted after the timer expires and after a join is received.
FwdCacheTimeout	Indicates the PIM forward cache expiry value in seconds. This value is used in aging PIM mroutes.
RegSupprTimeout	Indicates the register-suppression timer in seconds.
UniRouteChangeTimeout	Indicates the frequency at which the RTM is polled for routing information updates.

Field	Description
JoinPruneInt	Indicates the join pruning interval in seconds.

Configuring PIM on a VRF

Configure PIM to create or remove an instance of the multicast routing protocol on a Virtual Router Forwarding (VRF) instance. Use the VRF Lite feature with multicast routing protocols to create multiple virtual multicast routers. By default, PIM does not run on a VRF.

Prerequisites

- Multicast virtualization works on an R, RS, or 8800 module port.
- Multicast virtualization requires an 8692 SF/CPU with SuperMezz or 8895 SF/CPU.
- You must configure virtualization for the unicast routing system.
- You must log on to VRF Router Configuration mode.

Procedure steps

1. Create a PIM instance on a VRF:

   ```
   ip pim
   ```

2. Enable the PIM instance on a VRF:

   ```
   ip pim enable
   ```

3. Disable a PIM instance on a VRF:

   ```
   no ip pim enable
   ```

4. Remove a PIM instance from a VRF:

   ```
   no ip pim
   ```

5. Configure the remaining PIM parameters, as required.

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
bootstrap-period *<5-32757>*	Specifies the interval (in seconds) that the elected bootstrap router (BSR) waits between originating bootstrap messages. • *integer* is an integer in the range of 5–32757. The default is 60. To set this option to the default value, use the **default** operator with the command.
disc-data-timeout *<5–65535>*	Specifies how long (in seconds) to discard data until the join is received from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the the timer expires or a join is received. • *integer* is an integer in the range of 5–65535. The default is 60. To set this option to the default value, use the **default** operator with the command.
enable	Globally activates PIM on the switch. To set this option to the default value, use the **default** operator with the command. The default is disabled.
fast-joinprune	Enables the fast join prune interval. To set this option to the default value, use the **default** operator with the command. The default is no disabled.
fwd-cache-timeout *<10–8640>*	Specifies the forward cache timeout value. • *integer* is an integer in the range of 10–86400. The default is 210. **Important:** When you configure one of the timers, activity-chk-interval or fwd-cache-timeout, with a nondefault value, you cannot configure the other timer. To set this option to the default value, use the **default** operator with the command.
joinprune-interval *<1–18724>*	Specifies how long to wait (in seconds) before the PIM router sends out the next join/prune message to its upstream neighbors. • *integer* is an integer in the range of 1–18724. The default is 60. To set this option to the default value, use the **default** operator with the command.
mode *<sparse|ssm>*	Configures the mode of this interface globally. After you change from one mode to another, an information message appears to remind you that traffic does not stop

Variable	Value
	immediately. To set this option to the default value, use the **default** operator with the command. The default is sparse.
register-suppression-timeout <*6–65535*>	Specifies how long (in seconds) the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. • *integer* is an integer in the range of 6–65535. The default is 60. To set this option to the default value, use the **default** operator with the command.
rp-c-adv-timeout <*5–26214*>	Specifies how often (in seconds) a router configured as a candidate RP (C-RP) sends C-RP advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected BSR. • *integer* is an integer in the range of 5–26214. The default is 60. To set this option to the default value, use the **default** operator with the command.
rp-candidate group <A.B.C.D> <mask address> rp <A.B.C.D>	Adds or deletes candidate RP entries. Use the **no** operator to later remove this configuration.
static-rp <A.B.C.D/X> rp <A.B.C.D>	Adds static RP entries and activates static RP.
static-rp specific-route	With static RP enabled, if the route to the RP is removed, the Avaya Ethernet Routing Switch 8800/8600 can fail over to an alternate static RP. However, if a default route exists in the routing table, that default route still appears as an active route to the failed RP. In this case, the switch does not fail over to the alternate RP. A similar situation exists with SMLT-based configurations, where an internal-only default static route is used during IST failover and recovery. In this case, the internal default route appears as an active route to the failed RP, and therefore does not failover to the alternate RP. To resolve these situations, you can configure the lookup for static RP to be chosen from the specific route rather than the best route. In this case, when the route to the active RP fails, the switch no longer interprets the default route as a valid route for RP purposes, and therefore fails over to the alternate RP.
unicast-route-change-timeout <*2–65535*>	Specifies how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates for PIM.

Variable	Value
	Important: Lowering this value increases how often the switch polls the RTM. This can affect the performance of the switch, especially when a high volume of traffic flows through the switch. • *integer* is an integer in the range of 2–65535. The default is 5. To set this option to the default value, use the **`default`** operator with the command.
virtual-neighbor <A.B.C.D> <A.B.C.D>	Adds a virtual neighbor to an interface globally. A.B.C.D represents the IP addresses of the interface and the virtual neighbor. Use the **`no`** operator to later remove this configuration.

Configuring PIM on a VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable PIM globally before you configure PIM on a VLAN.
- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Enable PIM on a VLAN:

   ```
   ip pim enable
   ```

2. Configure the remaining parameters as required.
3. Verify your configuration by displaying information about the PIM-SM interface setup for VLANs:

   ```
   show ip pim interface vlan [vrf Word<0-16>] [vrfids Word<0-255>]
   ```

Variable definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
active	Enables PIM and configures the interface type to active.
bsr-candidate preference <pref value>	Enables the BSR candidate on a specific port. The preference value ranges from 0–255. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
enable	Enables PIM on the local switch interface. To set this option to the default value, use the **default** operator with the command. The default is disabled.
interface-type <active| passive>	Specifies whether the selected interface is active or passive. You can change the state of a PIM interface after you create the interface but only if you disable PIM on the interface. An active interface accepts PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful when a high number of PIM interfaces exist and connect to end users, not to other switches. To set this option to the default value, use the **default** operator with the command. The default is active.
joinprune-interval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default is 60 seconds. To set this option to the default value, use the **default** operator with the command. The range is 1-18724.
passive	Enables PIM and configures the interface type to passive.
query-interval <seconds>	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default is 30 seconds with the range of 0–18724. To set this option to the default value, use the **default** operator with the command.

Job aid

The following table shows the field descriptions for this command.

Field	Description
VLAN-ID	Identifies the VLAN.
PIM-ENABLE	Identifies the state of PIM on the VLAN.
MODE	Identifies the configured mode of this VLAN. The valid modes are SSM and Sparse.
HELLOINT	Indicates how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Indicates how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default join/prune interval is 60 seconds.
CBSR PREF	Indicates the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
INTF TYPE	Indicates whether the PIM interface is active or passive.

Configuring PIM on an Ethernet port

Configure PIM for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable PIM globally before you configure it on an interface.
- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.
- The interface uses a valid IP address.

Procedure steps

1. Enable PIM on an Ethernet port:

   ```
   ip pim enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
active	Enables PIM and configures the interface type to active.
bsr-candidate preference <pref value>	Enables the BSR candidate on a specific port. The preference value ranges from 0–255. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
enable	Enables PIM on the local switch interface. To set this option to the default value, use the **default** operator with the command. The default is disabled.
interface-type <active| passive>	Specifies whether the selected interface is active or passive. You can change the state of a PIM interface after you create the interface but only if you disable PIM on the interface. An active interface accepts PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful when a high number of PIM interfaces exist and connect to end users, not to other switches. To set this option to the default value, use the **default** operator with the command. The default is active.
joinprune-interval <*seconds*>	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default is 60 seconds. To set this option to the default value, use the **default** operator with the command. The range is 1-18724.
passive	Enables PIM and configures the interface type to passive.
query-interval <seconds>	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default is 30 seconds with the range of 0–18724. To set this option to the default value, use the **default** operator with the command.

Configuring SSM globally

Configure SSM to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction,

SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces running PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

You configure PIM on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP and OSPF, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — OSPF and RIP,* (NN46205-522).

  PIM requires a unicast protocol to forward multicast traffic within the network when it performs the Reverse Path Forwarding (RFP) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that PIM-enabled routers use to communicate. The unicast routing table must contain a route to every multicast source in the network as well as routes to PIM entities like the RPs and BSR.
- Enable PIM globally.
- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Configure PIM-SSM:

```
ip pim mode ssm
```

Configuring DVMRP globally

Configure the Distance Vector Multicast Routing Protocol (DVMRP) globally to enable or disable DVMRP and change default global parameters.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Configure DVMRP globally:

   ```
   ip dvmrp
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`ip dvmrp`** command.

Variable	Value
fwd-cache-timeout <integer>	Configures the forward cache timeout (in seconds). • *integer* is the range of 10–86400 seconds. The default value is 300 seconds. To set this option to the default value, use the **`default`** operator with the command.
generate-log	Enables or disables the DVMRP log. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
generate-trap	Enables or disables the DVMRP trap. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
leaf-timeout <integer>	Configures the length of time (in seconds) the router waits for a response from a neighbor before considering the attached network as a leaf network. • *integer* is the range of 25–4000 seconds. The default value is 125 seconds. To set this option to the default value, use the **`default`** operator with the command.

Variable	Value
nbr-probe-interval <integer>	Configures the time interval (in seconds) for the DVMRP router to send a neighbor probe message on its interface. • *integer* is the range of 5–30 seconds. The default value is 10 seconds. To set this option to the default value, use the **default** operator with the command.
nbr-timeout <*integer*>	Configures the length of time (in seconds) the router waits to receive a report from a neighbor before considering the connection inactive. • *integer* is the range of 35–8000 seconds. The default value is 35 seconds. To set this option to the default value, use the **default** operator with the command.
output-report-delay <integer>	Configures the time interval (in seconds) between DVMRP router update messages. • *integer* is the range of 10–2000 seconds. The default value is 60 seconds. To set this option to the default value, use the **default** operator with the command.
prune-resend	Sends prune messages every 3 minutes, to address the link failures at remote upstream switches. The feature is disabled by default. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command.
route-discard-timeout <*integer*>	Configures the route discard timeout (in seconds). • *integer* is the range of 25–8000. The default value is 260 seconds. To set this option to the default value, use the **default** operator with the command.
route-expiration-timeout <*integer*>	Configures the route expiration timeout (in seconds). • *integer* is the range of 20–4000 seconds. The default value is 140 seconds. To set this option to the default value, use the **default** operator with the command.

Variable	Value
route-switch-timeout <*integer*>	Configures the route switch timeout (in seconds). • *integer* is the range of 20–2000. The default value is 140 seconds. To set this option to the default value, use the **`default`** operator with the command.
show-next-hop-table	Enables or disables showing information about the DVMRP next hops. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
triggered-update-interval <*integer*>	Configures the time interval (in seconds) between triggered update messages sent after routing information changes. • *integer* is the range of 5–1000 seconds. The default value is 5 seconds. To set this option to the default value, use the **`default`** operator with the command.

Configuring DVMRP on a VLAN

Configure DVMRP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on a VLAN.
- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Enable DVMRP on a VLAN:

   ```
   ip dvmrp enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
active	Enables DVMRP on a specific interface with a specific type. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is active.
advertise-self	Enables or disables the advertisement of local routes for the selected interface to other switches in the network. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to advertise.
default-listen	Learns the default route over the specified interface if this feature is enabled on the interface. The default setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to listen.
default-supply	Generates and advertises the default route if this feature is enabled on the interface. The default setting is disable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to not supply.
default-supply-metric <*cost*>	Advertises the specified metric over the interface if you configured the interface to supply the default route. The range is 1–31 hops. The default setting is 1 hop. To set this option to the default value, use the **default** operator with the command.
enable	Enables DVMRP on the local router interface. To set this option to the default value, use the **default** operator with the command. The default is disabled.
in-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Variable	Value
interface-type active	Configures an interface as active. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is active.
metric <*cost*>	Configures the cost metric (maximum number of hops) for the router interface. • *cost* is the range of 1–31. To set this option to the default value, use the **default** operator with the command. The default is 1.
out-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Configuring DVMRP on Ethernet ports

Configure DVMRP for each interface to enable the interface to perform multicasting operations.

Prerequisites

- You must enable DVMRP globally before you enable it on an Ethernet port.
- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.
- The interface uses a valid IP address.

Procedure steps

1. Enable DVMRP on an Ethernet port:

   ```
   ip dvmrp enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
active	Enables DVMRP on a specific interface with a specific type. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is active.
advertise-self	Enables or disables the advertisement of local routes for the selected interface to other switches in the network. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to advertise.
default-listen	Learns the default route over the specified interface if this feature is enabled on the interface. The default setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to listen.
default-supply	Generates and advertises the default route if this feature enabled on the interface. The default setting is disable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to not supply.
default-supply-metric <*cost*>	Advertises the specified metric over the interface if you configured the interface to supply the default route. The range is 1–31 hops. The default setting is 1 hop. To set this option to the default value, use the **default** operator with the command.
enable	Enables DVMRP on the local router interface. To set this option to the default value, use the **default** operator with the command. The default is disabled.
in-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Variable	Value
interface-type active	Configures an interface as active. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is active.
metric <*cost*>	Configures the cost metric (maximum number of hops) for the router interface. • *cost* is the range of 1–31. To set this option to the default value, use the **default** operator with the command. The default is 1.
out-policy <*policy_name*>	Applies a DVMRP accept policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Configuring IGMP on a VLAN

Configure IGMP for each interface to change default multicasting operations. You can specify the version of IGMP and backward compatibility operations.

Prerequisites

- You must globally enable either PIM or DVMRP.
- You must enable either PIM or DVMRP on the VLAN.
- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

Configure IGMP on a VLAN:

```
ip igmp
```

Variable definitions

Use the data in the following table to use the **ip igmp** command.

Variable	Value
access-list	Displays the IP Multicast access group list.
compatibility-mode	Enables v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command: **default ip igmp compatibility-mode**, or use the no option to disable compatibility mode:**no ip igmp compatibility-mode**
dynamic-downgrade-version	Configures if the Avaya Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The switch downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command: **default ip igmp dynamic-downgrade-version** or use the no option to disable downgrade: **no ip igmp dynamic-downgrade-version**
igap	Enables the IGAP feature. To set this option to the default value, use the **default** operator with the command. The default is disabled.
immediate-leave	Enables fast leave mode. Prevents a port from receiving a leave message from a member of a group. Normal IGMP behavior is skipped. To set this option to the default value, use the **default** operator with the command. The default is disabled.
immediate-leave-members	Adds fast leave members for the VLAN.
last-member-query-interval <1/10_seconds>	The maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1.

Variable	Value
	Decreasing the value reduces the time to detect the loss of the last member of a group. • *1/10_seconds* is the range from 0–255, and the default is 10 tenths of a second. Avaya recommends that you configure this value between 3–10 (equal to 0.3–1.0 seconds).
mrdisc	Configures multicast router discovery parameters. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
mrouter <ports>	Adds multicast router ports.
proxy	Enables the proxy-snoop option globally for the VLAN. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
query-interval <seconds>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. • *seconds* is the range from 0–65535. The default value is 125 seconds.
query-max-response <1/10_seconds>	The maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. • *1/10_seconds* is an integer value with a range of 0–255, and the default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robust-value <integer>	Configures the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If a network is expected to lose query packets, increase the robustness value. • *integer* is an integer value with a range of 2–255 seconds. The default value is 2. The default value of 2 means that one query for each query interval is dropped without the querier aging out.
router-alert	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the

Variable	Value
	router processes IGMP packets regardless of whether the router alert IP option is set. **Important:** To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
snooping	Enables the snoop option for the VLAN. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
ssm-snoop	Enables support for PIM-SSM on the snoop interface. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
static-group	Displays the IP Multicast static parameters.
stream-limit	Enables the stream-limit feature.
stream-limit-group	Enables and configures stream-limit member features.
stream-limit-max-streams	Sets the maximum number of streams allowed on an interface.
version <integer>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. • *integer* is an integer value with a value of 1, 2, or 3. The default value is 2 (IGMPv2).

Example of configuring IGMP on a VLAN

Procedure steps

1. Enter VLAN Interface Configuration Mode for VLAN 1.

   ```
   interface vlan 1
   ```

2. Set the last member query interval to 15 tenths of a second (equal to 1.5 seconds).

```
ip igmp last-member-query-inte rval 15
```

3. Set the query interval to 100 seconds.

```
ip igmp query-interval 100
```

4. Set the query maximum response time to 15 tenths of a second (equal to 1.5 seconds).

```
ip igmp query-max-response 50
```

5. Set the robustness variable to 4.

```
ip igmp robust-value 4
```

6. Enable IGMPv3.

```
ip igmp version 3
```

7. Enable proxy snoop for the VLAN.

```
ip igmp proxy enable
```

8. Enable snoop for the VLAN.

```
ip igmp snooping enable
```

9. Enable support for SSM on the snoop interface.

```
ip igmp ssm-snoop enable
```

10. Enable the fast leave option.

```
ip igmp immediate-leave enable
```

Configuring IGMP Ethernet ports

Configure IGMP for each interface to change default multicasting operations. You can specify the version of IGMP and backward compatibility operations.

Prerequisites

- You must globally enable either PIM or DVMRP.
- You must enable either PIM or DVMRP on the port.
- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

Configure IGMP on Ethernet ports:

```
ip igmp
```

Variable definitions

Use the data in the following table to use the **ip igmp** command.

Variable	Value
access-list	Displays the IP Multicast access group list parameters.
compatibility-mode	Enables v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command: **default ip igmp compatibility-mode**, or use the no option to disable compatibility mode: **no ip igmp compatibility-mode**
dynamic-downgrade-version	Configures if the Ethernet Routing Switch 8800/8600 downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The switch downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command: **default ip igmp dynamic-downgrade-version**, or use the no option to disable downgrade: **no ip igmp dynamic-downgrade-version**
immediate-leave	Enables the fast leave option on the interface. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
last-member-query-interval <1/10_seconds>	Configures the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query

Variable	Value
	messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. • *1/10_seconds* is an integer in the range from 0–255 and the default is 10 tenths of a second. Avaya recommends that you configure this value between 3–10 (equal to 0.3–1.0 seconds). To set this option to the default value, use the **`default`** operator with the command.
port {slot/port[-slot/port][,...]}	Configures IGMP for a specific port.
query-interval <seconds>	Configures the frequency (in seconds) at which the interface transmits host query packets. • *seconds* is an integer in the range from 0–65535 with a default of 125. To set this option to the default value, use the **`default`** operator with the command.
query-max-response <1/10_seconds>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. • *1/10_seconds* is an integer value with a range of 0–255, and the default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval. To set this option to the default value, use the **`default`** operator with the command.
robust-value <integer>	Configures the expected packet loss of a network. • *integer* is an integer value with a range of 2–255 seconds. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss. To set this option to the default value, use the **`default`** operator with the command.
router-alert	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the

Variable	Value
	router processes IGMP packets regardless of whether the router alert IP option is set. **Important:** To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command.
stream-limit	Enables the stream-limit feature.
stream-limit-max-streams	Sets the maximum number of streams allowed on an interface.
version <integer>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. • *integer* is an integer value with a value of 1, 2, or 3. The default value is 2 (IGMPv2).

Deleting a single IP Multicast record

Use this procedure to delete a single IP Multicast record.

Prerequisites

- You must log on to the PrivExec Configuration mode in the ACLI.

Procedure steps

Delete a multicast record:

```
clear ip mroute group <ipaddr> vlan <vid> [source <ipaddr>]
```

Variable definitions

Use the data in the following table to use the **clear ip mroute group** command.

Variable	Value
group <ipaddr>	Specifies the multicast group address.
vlan <vid>	Specifies the VLAN ID.
[source <ipaddr>]	Specifies the multicast source address.

Chapter 8: DVMRP configuration using Enterprise Device Manager

Routers use the Distance Vector Multicast Routing Protocol (DVMRP) to exchange multicast routing information.

Prerequisites to DVMRP configuration

- Configure an IP interface.
- Disable Protocol Independent Multicast-Sparse Mode (PIM-SM) on the interface where you want to configure DVMRP because you cannot configure PIM-SM and DVMRP on the same interface.

  Important:

  Avaya recommends that you do not change the configuration from PIM to DVMRP, or from DVMRP to PIM, while multicast traffic is flowing on the network.

  A switch can use a mix of DVMRP and PIM-SM interfaces if it is configured as a multicast border router (MBR). You can configure an interface with only one multicast routing protocol at a time (DVMRP or PIM-SM).
- Enable DVMRP globally.

Important:

For DVMRP scaled configurations with more than one thousand streams, to avoid multicast traffic loss, increase routing protocol timeouts, for example, dead interval for OSPF.

Navigation

- Editing DVMRP interface parameters on page 180
- Editing DVMRP interface advance parameters on page 181
- Viewing DVMRP neighbor parameters on page 182
- Viewing DVMRP learned routes on page 183

- Viewing DVMRP next-hop information on page 184
- Applying the default route policy to an interface on page 185
- Applying the default route policy to a VLAN on page 186
- Applying the default route policy to a port on page 187
- Creating a DVMRP announce policy on page 187
- Applying a DVMRP announce policy to an interface on page 190
- Applying a DVMRP announce policy to a VLAN on page 191
- Applying a DVMRP announce policy to a port on page 192
- Creating a DVMRP accept policy on page 192
- Applying a DVMRP accept policy to an interface on page 194
- Applying a DVMRP accept policy to a VLAN on page 194
- Applying a DVMRP accept policy to a port on page 195
- Applying the advertisement of local networks policy over an interface on page 195
- Applying the advertisement of local networks policy over a VLAN on page 196
- Applying the advertisement of local networks policy over a port on page 196
- Configuring an active or passive interface type on page 197
- Configuring an active or passive VLAN type on page 197
- Configuring an active or passive port type on page 198

Editing DVMRP interface parameters

Edit DVMRP interface parameters to customize DVMRP at the interface level.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Interfaces** tab.
4. In the **Metric** field, change the metric value.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Interfaces tab.

Variable	Value
IfIndex	Shows the DVMRP interface, slot or port number, or VLAN identification.
LocalAddress	Shows the IP address of the DVMRP router interface.
Metric	Configures the distance metric for this interface; used to calculate distance vectors. The range is from 1–31. The default value is 1, which means local delivery only.
OperState	Shows the current operational state of this DVMRP interface (up or down).

Editing DVMRP interface advance parameters

Edit advance parameters to configure an interface to listen for a default route or supply a default route. You can also configure an interface to specify the metric (cost) of the default route, which is advertised if the interface is configured to supply a default route.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Interface Advance** tab.
4. Change the required values.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Interface Advance tab.

Variable	Value
IfIndex	Provides the DVMRP interface, VLAN, or slot or port number identification.

Variable	Value
LocalAddr	Provides the IP address of the DVMRP router interface.
Enable	Enables (true) or disables (false) DVMRP on the interface.
Metric	Specifies the distance metric for this interface; used to calculate distance vectors. The range is from 1–31 hops. The default is 1.
InterfaceType	Configures the interface type as passive or active. The default is active.
DefaultListen	Configures the interface to listen (true) or not listen (false) for the default route. The default is true, which indicates that the interface listens to the default route.
DefaultSupply	Configures the interface to supply (true) or not supply (false) only the default route. The default is false, which indicates not to supply a default route on that interface.
DefaultRouteMetric	Configures the metric (number of hops for DVMRP) of the default route. The range is from 1–31 hops. The default is 1.
AdvertiseSelf	Configures the interface to advertise (true) or not advertise (false) its local route to neighbors. The default value is true.
InPolicy	Specifies the name of the DVMRP accept policy applied to the interface.
OutPolicy	Specifies the name of the DVMRP announce policy applied to the interface.

Viewing DVMRP neighbor parameters

View the DVMRP neighbor parameters for troubleshooting purposes.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Neighbors** tab.

   The DVMRP parameters appears.

Variable definitions

Use the data in the following table to use the Neighbors tab.

Variable	Value
IfIndex	Shows the DVMRP slot or port number or the virtual interface (VLAN) used to reach this DVMRP neighbor.
Address	Shows the IP address of the DVMRP neighbor where this entry contains information.
ExpiryTime	Shows the time remaining before this DVMRP neighbor is aged out.
GenerationId	Shows the generation ID number for the neighboring router.
MajorVersion	Shows the major DVMRP version number for the neighboring router.
MinorVersion	Shows the minor DVMRP version number for the neighboring router.
Capabilities	Shows the capabilities for the neighboring router. The probe flag is 1 byte long with the lower 4 bits containing the following information: • The leaf bit (0) indicates that the neighbor uses only one interface with neighbors. • The prune bit (1) indicates that the neighbor supports pruning. • The generationID bit (2) indicates that the neighbor sends its generation ID in probe messages. • The mtrace bit (3) indicates that the neighbor can handle mtrace requests.
State	Shows the state of neighbor adjacency: • oneway—The switch sees a packet from the neighbor but no adjacency is established. • active—Adjacency exists in both directions. • ignoring • down—The interface is not enabled.

Viewing DVMRP learned routes

View the DVMRP learned routing table for troubleshooting purposes.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Routes** tab.

   DVMRP learned routing table appears.

Variable definitions

Use the data in the following table to use the Routes tab.

Variable	Value
Source	Shows the network address that, when combined with the corresponding route SourceMask value, identifies the sources containing multicast routing information.
SourceMask	Shows the network mask that, when combined with the corresponding route Source value, identifies the sources containing multicast routing information.
UpstreamNeighbor	Shows the address of the upstream neighbor from which the IP datagrams from these sources are received. The address is 0.0.0.0 if the network is local.
Interface	Shows the DVMRP interface, slot or port number, or VLAN ID where the IP datagrams sent by the upstream sources are received.
Metric	Shows the distance in hops to the source subnet. The range is from 1–32.
ExpiryTime	Shows the amount of time (in seconds) remaining before this entry is aged out.

Viewing DVMRP next-hop information

View information about the DVMRP next hops on outgoing interfaces for routing IP multicast datagrams. Showing the next-hop table is disabled by default. This avoids using the large amount of memory required for these tables in a scaled multicast environment with a large number of VLANs.

Prerequisites

- Before you can show DVMRP next hops, use the CLI to enable showing the next-hop table.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Next Hops** tab.

   DVMRP next-hop information appears.

Variable definitions

Use the data in the following table to use the Next Hops tab.

Variable	Value
Source	Shows the network address that, when combined with the corresponding next-hop SourceMask value, identifies the source of the next hop on an outgoing interface.
SourceMask	Shows the network mask that, when combined with the corresponding next-hop Source value, identifies the source of the next hop on an outgoing interface.
IfIndex	Shows the DVMRP interface, slot or port number, or VLAN ID for the outgoing interface for the next hop.
Type	Shows the type of next hop. The type is 0, or leaf, if no downstream dependent neighbors exist on the outgoing virtual interface. Otherwise, the type is branch.

Applying the default route policy to an interface

You can configure an interface to listen for a default route or supply a default route. You can also configure an interface to specify the metric (cost) of the default route, which is advertised if you configure the interface to supply a default route.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Interface Advance** tab.
4. Set the **DefaultListen** field for the interface you want to modify to **true** if you want to listen for a default route.
5. Set the **DefaultSupply** field for the interface you want to modify to **true** if you want the interface to supply only a default route,
6. Type the number of hops for DVMRP in the **DefaultRouteMetric** field to set the metric (cost) of the default route used when this switch advertises this default route; the range is 1 to 31.
7. Click **Apply** to save the new configuration.

Applying the default route policy to a VLAN

You can configure a VLAN to listen for a default route or supply a default route. You can also configure a VLAN to specify the metric (cost) of the default route, which is advertised if you configure the VLAN to supply a default route.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN ID that you want to configure.
4. Click **IP** from the menu bar.
5. Click the **DVMRP** tab.
6. To set the VLAN to listen for a default route, select the **DefaultListen** check box.
7. To set the VLAN to supply only the default route, select the **DefaultSupply** check box.
8. To set the metric (cost) of the default route used when this switch advertises this default route, type the number of hops for DVMRP in the **DefaultRouteMetric** field; the range is 1 to 31.
9. Click **Apply** to save the new configuration.

Applying the default route policy to a port

You can configure a port to listen for a default route or supply a default route. You can also configure a port to specify the metric (cost) of the default route, which is advertised if you configure the port to supply a default route.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DVMRP** tab.
5. To set the port to listen for a default route, select the **DefaultListen** check box.
6. To set the port to supply only the default route, select the **DefaultSupply** check box.
7. To set the metric (cost) of the default route used when this switch advertises this default route, type the number of hops for DVMRP in the **DefaultRouteMetric** field; the range is 1 to 31.
8. Click **Apply** to save the new configuration.

Creating a DVMRP announce policy

Create one or more IP prefix lists and apply those lists to an IP route policy. A prefix list with a 32-bit mask is equivalent to an address. You can use a prefix list with a mask less than 32 bits as a network. If you configure the MaskLenFrom field less than the MaskLenUpto field, you can also use it as a range.

After you create the announce policy, apply the policy to an interface, VLAN, or port.

Prerequisites

- Create an IP prefix list to use in the announce policy.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click the **Route Policy** tab.
4. Click **Insert**.

   The Insert Route Policy dialog box appears.

   Important:
   Not all the fields in the Policy, Insert Route Policy dialog box apply to the process of creating a DVMRP policy.
5. In the **Id** field, type number for the policy
6. In the **SequenceNumber** field, type a number for the policy (the range is 1–65535).
7. In the **Name** field, type a name for the policy.
8. Select the **Enable** check box to enable the policy.
9. In the **Mode** filed, click **permit** or **deny**, to permit or deny the policy.
10. In the **MatchProtocol** field, select **DVMRP**.
11. In the **MatchNetwork** box, click **[...]** button.

    The MatchNetwork dialog box appears.
12. Type the applicable names of the prefix lists you created and click **Ok**.
13. In the **MatchIPRouteSource** box, click **[...]** button.

    The MatchIPRouteSource dialog box appears.
14. Type the applicable names of the prefix lists you created and click **Ok**.
15. In the **MatchExtendedPref**, type a value to associate extended ACL (access control list) with the route policy.
16. In the **MatchNextHop** box, click **[...]** button.

    The MatchNextHop dialog box appears and click.
17. In the **MatchRouteType** field, select a specific route type to match.
18. In the **MatchMetric** field, type a value (the range is 0–65535).

19. In the **MatchVrf** field, click **[..]** button.

    MatchVrf dialog box appears.
20. Select a VRF and click **Ok**.
21. In the **NssaPbit** box, select the **Enable** check box.
22. In the **SetRoutePreference** box, type the value to be assigned to the routes.
23. In the **SetMetrictypeInternal**, type a value to set the MED value for routes.
24. In the **SetMetric** box, type a value (the range is 0–65535).
25. In the **SetInjectNetList** field, click **[...]** button. SetInjectNetList dialog box appears.
26. Select a route and click **Ok**.
27. Click **Insert**.
28. Click **Apply** to save the new configuration.

Variable definitions

Use the data in the following table to configure the Policy, Insert Route Policy dialog box.

Variable	Value
Id	The ID of an entry in the prefix list table.
SequenceNumber	A second index used to identify a specific policy within a route policy group.
Name	The name of the route policy.
Enable	Enables the policy.
Mode	Permits or denies the policy.
MatchProtocol	Selects the appropriate protocol. If you configure this variable, matches the protocol through which the route is learned.
MatchNetwork	Matches the destination network against the contents of the specified prefix list.
MatchIpRouteSource	Matches the previous hop IP addresses for DVMRP routes against the contents of the specified prefix list, if configured. Click the ellipsis button and choose from the list in the Match Route Source dialog box. You can select up to four entries. To clear an entry, use the ALT key. **Important:** You can also change this field in the Route Policy tab of the Policy dialog box.

Variable	Value
MatchNextHop	Matches the previous hop IP address of the route against the contents for the specified prefix list. This field applies only to nonlocal routes, if configured. Click the ellipsis button and choose from the list in the Match Next Hop dialog box. You can select up to four entries. To clear an entry, use the ALT key.
MatchInterface	If configured, the switch matches the IP address of the interface by the RIP route that was learned against the contents of the specified prefix list. This field is used only for RIP routes and it is ignored for all other route types. Click the ellipsis button and choose from the list in the Match Interface dialog box. You can select up to four entries. To clear an entry, use the ALT key.
MatchRouteType	Configures a specific route type to match. This setting applies only to Open Shortest Path First (OSPF) routes. Externaltype1 and Externaltype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra and interarea routes.
MatchMetric	Matches the metric of the incoming advertisement or existing route against the specified value (1–65535). If 0, this field is ignored. The default is 0.
MatchAsPath	If configured, the switch matches the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified as-lists. This field is used only for BGP routes and ignored for all other route types.
SetMetric	If configured, the switch configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or default-import-metric is used. Use the no operator to later remove this configuration.

Applying a DVMRP announce policy to an interface

Apply an announce policy to an interface to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Select the **Interface Advance** tab.
4. Click the **OutPolicy** field.

   The OutPolicy the screen appears, displaying the policy names.
5. Select an interface.
6. Click **OK**.
7. Click **Apply**.

Applying a DVMRP announce policy to a VLAN

Apply an announce policy to a VLAN to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
3. Click **VLANs**.
4. Select a VLAN, and then click **IP**.
5. Click the **DVMRP** tab.
6. In **OutPolicy** box, click **[...]** button.

   The OutPolicy screen appears.
7. Select a policy name.
8. Click **OK**
9. Click **Apply**.

Applying a DVMRP announce policy to a port

Apply an announce policy to a port to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DVMRP** tab.
5. In **OutPolicy** box, click **[...]** button.
6. Select a policy name.
7. Click **OK**.
8. Click **Apply**.

Creating a DVMRP accept policy

Create one or more IP prefix lists and apply those lists to an IP route policy. A prefix list with a 32-bit mask is equivalent to an address. You can use a prefix list with a mask less than 32 bits as a network. If you configure the MaskLenFrom field less than the MaskLenUpto field, you can also use it as a range.

After you create the policy, you can apply it to an interface, VLAN, or port.

Prerequisites

- Configure an IP prefix list to use in the accept policy.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Policy**.

3. Click the **Route Policy** tab.
4. Click **Insert**.

   The Insert Prefix List dialog box appears.

   Important:

   Not all the fields in the Policy, Insert Route Policy dialog box apply to the process of creating a DVMRP policy.
5. In the **Id** field, type number for the policy
6. In the **SequenceNumber** field, type a number for the policy (the range is 1–65535).
7. In the **Name** field, type a name for the policy.
8. Select the **Enable** check box to enable the policy.
9. In the **MatchProtocol** field, select **DVMRP**.
10. In the **MatchNetwork** box, click **[...]** button.

    The MatchNetwork dialog box appears.
11. Type the applicable names of the prefix lists you created.
12. Click **Ok**.
13. In the **MatchIPRouteSource** box, click **[...]** button.

    The MatchIPRouteSource dialog box appears.
14. Type the applicable names of the prefix lists you created.
15. Click **Ok**.
16. In the **MatchNextHop** box, click **[...]** button.

    The MatchNextHop dialog box appears.
17. Type the applicable names of the prefix lists you created.
18. Click **Ok**.

    You can use a single list or several prefix lists. You can select up to four lists. To select the names of the prefix lists, click the ellipsis button to the right of the field, select the appropriate names from the dialog box, and then click OK. To select multiple names, use the CTRL key. To deselect an entry, use the ALT key.
19. In the **MatchMetric** field, type a value (the range is 0–65535).
20. In the **SetMetric** box, type a value (the range is 0–65535).
21. Click **Insert**.
22. From the **Policy** dialog box, click the **DVMRP In/Out Policy** tab.

23. Right-click in the **InPolicy** box of the DVMRP interface to which you want to apply the accept policy, and then select the appropriate policy name from the **PolicyName** dialog box.
24. Click **Apply** to save the new configuration.

Applying a DVMRP accept policy to an interface

Apply an accept policy to an interface to control the way DVMRP manages incoming routes advertisements.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Select the **Interface Advance** tab.
4. Click the **InPolicy** field for a selected interface.
5. Select a policy name, and then click **OK**.
6. Click **Apply**.

Applying a DVMRP accept policy to a VLAN

Apply an accept policy to a VLAN to control the way DVMRP manages incoming routes advertisements.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN, and then click **IP** from the menu bar.
4. Click the **DVMRP** tab.
5. Click the tab to the right of the **InPolicy** field.
6. Select a policy name, and then click **OK**.
7. Click **Apply**.

Applying a DVMRP accept policy to a port

Apply an accept policy to a port to control the way DVMRP manages incoming routes advertisements.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DVMRP** tab.
5. Click the tab to the right of the **InPolicy** field.
6. Select a policy name, and then click **OK**.
7. Click **Apply**.

Applying the advertisement of local networks policy over an interface

Apply the advertisement of local networks policy to advertise local networks over an interface.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **DVMRP**.
3. Click the **Interface Advance** tab.
4. To enable the advertisement of local networks policy for a selected interface, set the **AdvertiseSelf** field for the interface you want to modify to **true**.
5. Click **Apply** to save the new configuration.

Applying the advertisement of local networks policy over a VLAN

Apply the advertisement of local networks policy to advertise local networks over a VLAN.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN, and then click **IP** from the menu bar.
4. Click the **DVMRP** tab.
5. To enable the VLAN to advertise its local networks, select the **AdvertiseSelf** checkbox.
6. Click **Apply** to save the new configuration.

Applying the advertisement of local networks policy over a port

Apply the advertisement of local networks policy to advertise local networks over a port.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DVMRP** tab.
5. To enable the port to advertise its local networks, select the **AdvertiseSelf** check box.
6. Click **Apply** to save the new configuration.

Configuring an active or passive interface type

Configure an interface as active or passive.

Prerequisites

- The interface is disabled.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP,** .
2. Click **DVMRP**.
3. Click the **Interface Advance** tab.
4. Select an interface.
5. In the **Enable** field, select **false** to disable the interface.
6. Click the **InterfaceType** field and then from the list set the interface to **passive** or **active**.
7. Click the **Enable** field and then from the list set the interface to **true** or **false** option.
8. Click **Apply** to save the new configuration.

Configuring an active or passive VLAN type

Configure a VLAN as active or passive.

Prerequisites

- The interface is disabled.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
3. Click **VLANs**.
4. Select a VLAN, and then click **IP** from the menu bar.
5. Click the **DVMRP** tab.
6. Clear the **Enable** check box to disable the interface.
7. In the **InterfaceType** box, set the interface to **passive** or **active**.
8. Select the **Enable** check box to reenable the interface.
9. Click **Apply** to save the new configuration.

Configuring an active or passive port type

Configure a port as active or passive.

Prerequisites

- The interface is disabled.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DVMRP** tab.

5. Clear the **Enable** check box to disable the interface.
6. In the **InterfaceType** box, set the interface to **passive** or **active**.
7. Select the **Enable** check box to reenable the interface.
8. Click **Apply** to save the new configuration.

Chapter 9: DVMRP configuration using the CLI

Configure the Distance Vector Multicast Routing Protocol (DVMRP) to exchange multicast routing information between routers.

For more information about DVMRP configuration examples, see *Technical Configuration Guide for Passport 8600 Multicast DVMRP*. You can find this document at http://www.avaya.com/support.

Prerequisites to DVMRP configuration

- Enable DVMRP globally.
- Disable Protocol Independent Multicast-Sparse Mode (PIM-SM) on the interface where you want to configure DVMRP because you cannot configure PIM-SM and DVMRP on the same interface.

  Important:

  Avaya recommends that you do not change the configuration from PIM to DVMRP, or from DVMRP to PIM, while multicast traffic is flowing on the network.
- A switch can use a mix of DVMRP and PIM-SM interfaces if it is configured as a multicast border router (MBR).

Important:

For DVMRP scaled configurations with more than one thousand streams, to avoid multicast traffic loss, increase routing protocol timeouts, for example, dead interval for OSPF.

DVMRP configuration navigation

- Job aid on page 202
- Showing DVMRP next hops on page 204
- Applying the default route policy to an interface on page 205
- Applying the default route policy to a VLAN on page 207

- Applying the default route policy to a port on page 209
- Creating a DVMRP policy on page 211
- Applying a DVMRP announce policy to an interface on page 217
- Applying a DVMRP announce policy to a VLAN on page 218
- Applying a DVMRP announce policy to a port on page 219
- Applying a DVMRP accept policy to an interface on page 220
- Applying a DVMRP accept policy to a VLAN on page 221
- Applying a DVMRP accept policy to a port on page 222
- Applying the advertisement of local networks policy to an interface on page 223
- Applying the advertisement of local networks policy over a VLAN on page 224
- Applying the advertisement of local networks policy over a port on page 225
- Creating a passive interface on page 226
- Configuring an active or passive interface type on page 227
- Configuring an active or passive VLAN type on page 228
- Configuring an active or passive port type on page 229

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ip route-policy <policy_name> seq <seq_number>`	`info`
	`action <permit|deny>`
	`create`
	`delete`
	`enable`
	`disable`
	`match-as-path <as-list> [clear]`
	`match-community <community-list> [clear]`

Command	Parameter
	match-community-exact <enable|disable> [clear]
	match-interface <prefix-list> [clear]
	match-metric <metric> [clear]
	match-network <prefix-list> [clear]
	match-next-hop <prefix-list> [clear]
	match-protocol <protocol_name> [clear]
	match-route-src <prefix-list> [clear]
	match-route-type <route-type>
	match-tag <tag> [clear]
	name <policy_name>
	set-as-path <as-list-id> [clear]
	set-as-path-mode <tag| prepend> [clear]
	set-automatic-tag <enable| disable> [clear]
	set-community <community-list> [clear]
	set-community-mode <additive| none> [clear]
	set-injectlist <prefix-list> [clear]
	set-local-pref <pref-value> [clear]
	set-mask <ipaddr>
	set-metric <metric-value> [clear]
	set-metric-type <metric-type> [clear]

Command	Parameter
	set-nssa-pbit <enable| disable>
	set-next-hop <ipaddr> [clear]
	set-origin <origin> [clear]
	set-origin-egp-as <origin-egp-as> [clear]
	set-preference <pref-value> [clear]
	set-tag <tag> [clear]
	set-weight <weight> [clear]
show ip dvmrp next-hop	
show ip dvmrp show-all [file <value>]	

Showing DVMRP next hops

Showing the next-hop table is disabled by default. Disable this setting to avoid using the large amount of memory required for these tables in a scaled multicast environment with a large number of VLANs.

Procedure steps

1. Enable showing the next-hop table:

   ```
   config ip dvmrp show-next-hop-table enable
   ```

2. Display information about the DVMRP next hops:

   ```
   show ip dvmrp next-hop
   ```

Job aid

The following table shows the field descriptions for this command.

Field	Description
SOURCE	Indicates the network address that, when combined with the corresponding value of dvmrpRouteNextHopSourceMask, identifies the sources for a next hop on an outgoing interface.
MASK	Indicates the network mask that, when combined with the corresponding value of dvmrpRouteNextHopSource, identifies the sources for a next hop on an outgoing interface.
INTERFACE	Indicates the outgoing interface for this next hop.
TYPE	Displays leaf if no downstream dependent neighbors exist on the outgoing virtual interface. Otherwise, the type is branch.

Applying the default route policy to an interface

Apply the default route policy to an interface. With DVMRP routing policies, you can improve the management of the DVMRP routing tables by controlling how the routing table is populated and how Avaya Ethernet Routing Switch 8800/8600 devices exchange the routes.

Procedure steps

1. Apply the policy to an interface:

   ```
   config ip dvmrp interface <ipaddr> default-listen <enable|
   disable> default-supply <enable|disable> default-supply-
   metric <cost>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ip dvmrp interface`** command.

Variable	Value
advertise-self <enable|disable>	Advertises the local network if the setting is enable. The default is enable.
create <active|passive>	Enables DVMRP on a specific interface of a specific interface type. • *active* is an active interface type. • *passive* is a passive interface type.

Variable	Value
	The default is active.
default-listen <enable|disable>	Learns the default route over the specified interface if this feature is enabled on the interface. The options are enable and disable. The default setting is enable.
default-supply <enable|disable>	Generates and advertises the default route if this feature is enabled on the interface. The options are enable and disable. The default setting is disable.
default-supply-metric <*cost*>	Advertises the specified metric over the interface if you configured the interface to supply the default route. The range is 1–31 hops. The default setting is 1 hop.
disable	Disables DVMRP on a specific interface.
enable	Enables DVMRP on a specific interface
in-policy <*policy_name*>	Specifies the DVMRP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64.
info	Shows the current level parameter settings and next level directories.
interface-type <active|passive>	Specifies the interface type as active or passive. The default is active.
ipaddr	Indicates the IP address of the selected interface.
metric <*cost*>	Specifies the DVMRP route metric. • *cost* is the metric value with a range of 1–31. The default is 1.
out-policy <*policy_name*>	Specifies the DVMRP route out-policy. • *policy_name* is a policy name that uses a string length of 0–64.

Example of applying the default route policy to an interface

Procedure steps

1. Enable the switch to learn the default route over interface 100.100.100.2.

```
8610:5# config ip dvmrp interface 100.100.100.2 default-listen enable
```

2. Enable the switch to generate and advertise the default route over interface 100.100.100.6.

```
8610:5# config ip dvmrp interface 100.100.100.6 default-supply enable

8610:5/config/ip/dvmrp/interface/100.100.100.6# default-supply-metric 3
```

Applying the default route policy to a VLAN

Apply the default route policy to a VLAN. With DVMRP routing policies, you can improve the management of the DVMRP routing tables by controlling how the routing table is populated and how Avaya Ethernet Routing Switch 8800/8600 devices exchange the routes.

Procedure steps

1. Apply the default route policy:

```
config vlan <vid> ip dvmrp default-listen <enable|disable>
default-supply <enable|disable> default-supply-metric <cost>
```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip dvmrp`** command.

Variable	Value
advertise-self <enable|disable>	Advertises the local network if set to enable. The default is enable.
create <active|passive>	Enables DVMRP on a specific interface of a specific interface type. • **`active`** is an active interface type. • **`passive`** is a passive interface type.

Variable	Value
	The default is active.
default-listen <enable|disable>	Learns the default route over the specified VLAN if this feature is enabled on the VLAN. The options are enable and disable. The default setting is enable.
default-supply <enable|disable>	Generates and advertises only the default route if this feature is enabled on the VLAN. No other route is advertised to the neighbors on this VLAN. The options are enable and disable. The default setting is disable.
default-supply-metric <*cost*>	Advertises the specified metric over the VLAN if you configure the VLAN to supply the default route. The range is 1–31 hops. The default setting is 1 hop.
disable	Disables DVMRP on a specific interface.
enable	Enables DVMRP on a specific interface.
in-policy <*policy_name*>	Specifies the DVRMP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64.
info	Shows the current level parameter settings and next level directories.
interface-type <active|passive>	Specifies the interface type as active or passive. The default is active.
metric <*cost*>	Specifies the DVMRP route metric. • *cost* is the metric value with a range of 1–31. The default is 1.
out-policy <*policy_name*>	Specifies the DVRMP route out-policy. • *policy_name* is a policy name that uses a string length of 0–64.
vid	Specifies a VLAN ID from 1–4092.

Example of applying the default route policy to a VLAN

Procedure steps

Enable the switch to learn the default route over VLAN 100 and to disable the VLAN from advertising the default route.

```
8610:5# config vlan 100 ip dvmrp default-listen enable
8610:5/config/vlan/100/ip/dvmrp# default-supply disable
```

Applying the default route policy to a port

Apply the default route policy to a port. With DVMRP routing policies, you can improve the management of the DVMRP routing tables by controlling how the routing table is populated and how Avaya Ethernet Routing Switch 8800/8600 devices exchange the routes.

Procedure steps

1. Apply the policy to a port:

   ```
   config ethernet <ports> ip dvmrp default-listen <enable|
   disable> default-supply <enable|disable> default-supply-
   metric <cost>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ethernet ip dvmrp`** command.

Variable	Value
advertise-self <enable|disable>	Advertises the local network if set to enable. The default is enable.
create <active|passive>	Enables DVMRP on a specific interface of a specific interface type. • **`active`** is an active interface type. • **`passive`** is a passive interface type. The default is active.

Variable	Value
default-listen <enable|disable>	Learns the default route over the specified port if this feature is enabled on the port. The options are enable and disable. The default setting is enable.
default-supply <enable|disable>	Generates and advertises only the default route if this feature is enabled on the port. No other route is advertised to the neighbors on this port. The options are enable and disable. The default setting is disable.
default-supply-metric <*cost*>	Advertises the specified metric over the port if you configure the port to supply the default route. The range is 1–31 hops. The default setting is 1 hop.
disable	Disables DVMRP on a specific interface.
enable	Enables DVMRP on a specific interface.
in-policy <*policy_name*>	Specifies the DVRMP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64.
info	Displays current configuration information.
interface-type <active|passive>	Specifies the interface type as active or passive. The default is active.
metric <*cost*>	Specifies the DVMRP route metric. • *cost* is the metric value with a range of 1–31. The default is 1.
out-policy <*policy_name*>	Specifies the DVRMP route out-policy. • *policy_name* is a policy name that uses a string length of 0–64.
ports	Specifies the port using the convention {slot/port[-slot/port][, ...]}.

Example of applying the default route policy to a port

Procedure steps

Configure port 9 of the card in slot 1 to listen for the default route and to disable the port from advertising the default route.

```
8610:5# config ethernet 1/9 ip dvmrp default-listen enable
8610:5/config/ethernet/1/9/ip/dvmrp# default-supply disable
```

Creating a DVMRP policy

Before you can apply an announce or accept policy to an interface, VLAN, or port, you must first create and configure the policy.

Prerequisites

- You must configure an IP prefix list to use in the policy.

Procedure steps

1. Create a DVMRP policy:

   ```
   config ip route-policy <policy-name> seq <seq_number> create
   ```

2. Create an action for the policy:

   ```
   config ip route-policy <policy-name> seq <seq_number> action
   <permit|deny>
   ```

3. Create a match condition.
4. Enable the policy:

   ```
   config ip route-policy <policy-name> seq <seq_number> enable
   ```

5. Configure the remaining parameters as required.

Important:

Not all of the **`route-policy seq`** parameters apply to the process of creating a DVMRP policy. The following table describes the parameters that you must use to create the DVMRP

policy. For more information about the other parameters for this command, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing, NN46205-523*.

Variable definitions

Use the data in the following table to use the **`config ip route-policy seq`** command.

Variable	Value
action <permit|deny>	Specifies the action to take when a policy matches a specific route. Permit accepts the route and deny ignores the route.
create	Creates a route policy with a policy name and a sequence number. **Important:** When you create a route policy in the CLI, the system internally generates the ID using an automated algorithm. When you create a route policy in Enterprise Device Manager, you can manually assign the ID number.
delete	Deletes a route policy with a policy name and a sequence number.
disable	Disables a route policy with a policy name and a sequence number.
enable	Enables a route policy with a policy name and a sequence number.
info	Displays current configuration information about this policy sequence number.
match-as-path <*as-list*> [clear]	If configured, the switch matches the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified as-lists. This field is used only for BGP routes and ignored for all other route types. • *as-list* specifies the list IDs of up to four as-lists, separated by commas. • **`[clear]`** removes the configured value for **`match-as-path`**.
match-community <*community-list*> [clear]	If configured, the switch matches the community attribute of the BGP routes against the contents of the specified community-lists. This field is used only for BGP routes and ignored for all other route types.

Variable	Value
	• *community-list* specifies the list IDs of up to four defined community-lists, separated by commas. • **[clear]** removes the configured value for **match-community**.
match-community-exact <enable| disable> [clear]	When disabled, **match-community** results in a match when the community attribute of the BGP routes matches an entry of a community list specified in **match-community**. When enabled, **match-community** results in a match when the community attribute of the BGP routes matches all of the entries of all the community lists specified in **match-community**.
match-interface <*prefix-list*> [clear]	If configured, the switch matches the IP address of the interface from which the Routing Information Protocol (RIP) route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other route types. • *prefix-list* specifies the names of up to four defined prefix lists, separated by commas. • **[clear]** removes the configured value for **match-interface**.
match-metric <*metric*> [clear]	If configured, the switch matches the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored. • *metric* is 1–65535. The default is 0. • **[clear]** removes the configured value for **match-metric**.
match-network <*prefix-list*> [clear]	If configured, the switch matches the destination network against the contents of the specified prefix lists. • *prefix-list* specifies the names of up to four defined prefix lists, separated by commas. • **[clear]** removes the configured value for **match-network**.
match-next-hop <*prefix-list*> [clear]	If configured, matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.

Variable	Value
	• *prefix-list* specifies the names of up to four defined prefix lists, separated by a comma. • **[clear]** removes the configured value for **match-next-hop**.
match-protocol <*protocol_name*> [clear]	If configured, matches the route policy to the protocol from which the route is learned. This field is used only for RIP announcement purposes.
match-route-src <*prefix-list*> [clear]	If configured, matches the next-hop IP address for RIP, BGP, and DVMRP routes and advertising router IDs for Open Shortest Path First (OSPF) routes against the contents of the specified prefix list. This option is ignored for all other route types. • *prefix-list* specifies the names of up to four defined prefix lists, separated by a comma. • **[clear]** removes the configured value for **match-route-src**.
match-route-type <*route-type*>	Configures a specific route-type to match (applies only to OSPF routes). • *route-type* External-1 and External-2 specifies OSPF routes of the specified type only (other values are ignored).
match-tag <*tag*> [clear]	Specifies a list of tags that is used during the match criteria process. Contains one or more tag values. • *tag* is a value from 0–256. • **[clear]** removes the configured values for **match-tag**.
name <*policy_name*>	Renames a policy after its creation. This command changes the name field for all sequence numbers under the policy.
policy-name	Indicates the name of the specified policy, which is a string length from 1–64 characters.
seq_number	Indicates the number of the specified policy, which is a number from 1–65535.
set-as-path <*as-list-id*> [clear]	If configured, the switch adds the as-number of the as-list to the BGP routes that match this policy. • *as-list-id* specifies the list ID of up to four defined as-lists, separated by a comma. • **[clear]** removes the configured value for **set-as-path**.

Variable	Value
set-as-path-mode <tag|prepend> [clear]	**prepend** is the default configuration. The switch prepends the as-number of the as-list specified in set-as-path to the old as-path attribute of the BGP routes that match this policy.
set-automatic-tag <enable| disable> [clear]	Configures the tag automatically. This option is used for BGP routes only.
set-community <*community-list*> [clear]	If configured, the switch adds the community number of the community list to the BGP routes that match this policy. • *community-list* specifies the list ID of up to four defined community lists, separated by a comma. • **[clear]** removes the configured value for **set-community**.
set-community-mode <additive| none> [clear]	Configures the community mode. • **additive**—the switch prepends the community number of the community list specified in set-community to the old community path attribute of the BGP routes that match this policy. • **none**—the switch removes the community path attribute of the BGP routes that match this policy. • **[clear]** removes the configured value for **set-community-mode**.
set-injectlist <*prefix-list*> [clear]	If configured, the switch replaces the destination network of the route that matches this policy with contents of the specified prefix list. • *prefix-list* specifies one prefix list by name. • **[clear]** removes the configured value for **set-injectlist**.
set-local-pref <*pref-value*> [clear]	A value used during a route decision process in the BGP protocol. Applicable to BGP only.
set-mask <*ipaddr*>	If configured, the switch configures the mask of the route that matches this policy. This applies only to RIP accept policies. • *ipaddr* is a valid contiguous IP mask.
set-metric <*metric-value*> [clear]	If configured, the switch configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or default-import-metric is used.

Variable	Value
set-metric-type <*metric-type*> [clear]	If configured, sets the metric type for the routes to announce into the OSPF domain that matches this policy. The default is type 2. This variable is applicable only for OSPF announce policies.
set-nssa-pbit <enable|disable>	Configures the not-so-stubby-area (NSSA) translation P bit. This variable is applicable only to OSPF announce policies.
set-next-hop <*ipaddr*> [clear]	Specifies the IP address of the next-hop router. Ignored for DVMRP routes.
set-origin <*origin*> [clear]	If configured, the switch changes the origin path attribute of the BGP routes that match this policy to the specified value.
set-origin-egp-as <*origin-egp-as*> [clear]	Indicates the remote autonomous system number. This variable is applicable to BGP only.
set-preference <*pref-value*> [clear]	Setting the preference greater than zero specifies the route preference value to assign to the routes that match this policy. This applies to accept policies only. • *pref-value* configures the preference value from 0–255. The default is 0. If the default is configured, the global preference value is used. • **[clear]** removes the configured value for **set-preference**.
set-tag <*tag*> [clear]	Configures the tag of the destination routing protocol. If you do not specify a tag, the switch forwards the tag value in the source routing protocol. A value of zero indicates that this parameter is not set.
set-weight <*weight*> [clear]	The weight value for the routing table. For BGP, this value overrides the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. Used for BGP only. A value of zero indicates that this parameter is not set.

Example of creating a DVMRP announce policy

Procedure steps

Configure the policy.

```
ERS_8606:5# config ip route-policy policy1
ERS_8606:5/config/ip/route-policy/policy1# seq 1
ERS_8606:5/config/ip/route-policy/policy1/seq/1#create
ERS_8606:5/config/ip/route-policy/policy1/seq/1#action deny
```

```
ERS_8606:5/config/ip/route-policy/policy1/seq/1# enable
ERS_8606:5/config/ip/route-policy/policy1/seq/1# match-network prefix1
```

Applying a DVMRP announce policy to an interface

Apply a DVMRP announce policy to an interface to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or provide a level of security for the network.

Prerequisites

- You must create the announce policy before you can apply it to an interface.

Procedure steps

To apply the announce policy to an interface, enter:

```
config ip dvmrp interface <ipaddr> out-policy <policy_name>
```

To later delete the policy from the VLAN, ener the following:

```
config ip dvmrp interface <ipaddr> out-policy "".
```

Variable definitions

Use the data in the following table to use the **`config ip dvmrp interface out-policy`** command.

Variable	Value
""	Indicates the policy used in the current session. Not entering a name for the out-policy deletes the current policy. **Important:** If you delete an announce policy from an interface, VLAN, or port, you change the configuration; you do not delete the policy itself.
ipaddr	Specifies the IP address of the interface.

Variable	Value
policy_name	Specifies the name of the announce policy you create.

Example of applying a DVMRP announce policy to an interface

Procedure steps

Apply an announce policy named policy1 to interface 2.2.2.2.

```
ERS_8606:5# config ip dvmrp interface 2.2.2.2 out-policy policy1
```

Applying a DVMRP announce policy to a VLAN

Apply a DVMRP announce policy to a VLAN to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Procedure steps

Apply a DVMP announce policy:

```
config vlan <vid> ip dvmrp out-policy <policy-name>
```

To later delete the policy from the VLAN, use the following command: **`config vlan <vid> ip dvmrp out-policy ""`**.

Variable definitions

Use the data in the following table to use the **`config vlan ip dvmrp out-policy`** command.

Variable	Value
""	Indicates the policy used in the current session. Not entering a name for the out-policy deletes the current policy. **Important:** If you delete an announce policy from an interface, VLAN, or port, you change the

Variable	Value
	configuration; you do not delete the policy itself.
policy-name	Specifies the number of the VLAN.
vid	Specifies the name of the announce policy to apply to the VLAN.

Example of applying a DVMRP announce policy to a VLAN

Procedure steps

Apply announce policy policy1 to VLAN 5.

```
ERS-8606:5# config vlan vlan5 ip dvmrp out-policy policy1
```

Applying a DVMRP announce policy to a port

Apply a DVMRP announce policy to a port to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Procedure steps

Apply a DVMRP announce policy:

```
config ethernet <port> ip dvmrp out-policy <policy-name>
```

To later delete the policy from the port, use the following command:

```
config ethernet <port> ip dvmrp out-policy "".
```

Variable definitions

Use the data in the following table to use the **`config ether ip dvmrp out-policy`** command.

Variable	Value
""	Indicates the policy used in the current session. Not entering a name for the out-policy deletes the current policy. **Important:** If you delete an announce policy from an interface, VLAN, or port, you change the configuration; you do not delete the policy itself.
policy-name	Specifies the name of the policy to apply to the port.
port	Specifies the number of the port.

Example of applying a DVMRP announce policy to a port

Procedure steps

Apply policy1 to port 1/5.

```
ERS_8606:5# config ether 1/5 ip dvmrp out-policy policy1
```

Applying a DVMRP accept policy to an interface

Apply a DVMRP accept policy to an interface to control the way DVMRP manages incoming route advertisements.

Procedure steps

Apply an accept policy:

```
config ip dvmrp interface <ipaddr> in-policy <policy-name>
```

To later delete the policy from the interface, use the following command:

```
config ip dvmrp interface <ipaddr> in-policy "".
```

Variable definitions

Use the data in the following table to use the **config ip dvmrp interface in-policy** command.

Variable	Value
""	Indicates the policy used in the current session. If you do not enter a name, you delete the policy. **Important:** If you delete an accept policy from an interface, VLAN, or port, you change the configuration; you do not delete the policy itself.
ipaddr	Specifies the address of the interface.
policy name	Specifies the name of the accept policy you create.

Example of applying a DVMRP accept policy to an interface

Procedure steps

Apply the policy named policy2 to interface 3.3.3.3.

```
ERS_8606:5# config ip dvmrp interface 3.3.3.3 in-policy policy2
```

Applying a DVMRP accept policy to a VLAN

Apply a DVMRP accept policy to a VLAN to control the way DVMRP manages incoming route advertisements.

Procedure steps

Apply an accept policy:

```
config vlan <vid> ip dvmrp in-policy <policy-name>
```

To later delete the policy from the VLAN, use the following command:

```
config vlan <vid> ip dvmrp in-policy "".
```

Variable definitions

Use the data in the following table to use the **`config vlan ip dvmrp in-policy`** command.

Variable	Value
""	Indicates the policy used in the current session. If you do not enter a name, you delete the policy. **Important:** If you delete an accept policy from an interface, VLAN, or port, you change the configuration; you do not delete the policy itself.
policy name	Specifies the name of the accept policy you create.
vid	Specifies the number of the VLAN.

Example of applying a DVMRP accept policy to a VLAN

Procedure steps

Apply the policy.

```
ERS_8606:5# config vlan vlan3 ip dvmrp in-policy policy2
```

Applying a DVMRP accept policy to a port

Apply a DVMRP accept policy to a port to control the way DVMRP manages incoming route advertisements.

Procedure steps

Apply an accept policy:

```
config ether <port> ip dvmrp in-policy <policy-name>
```

To later delete the policy from the port, use the following command:

```
config ether <port> ip dvmrp in-policy "".
```

Variable definitions

Use the data in the following table to use the **`config ether ip dvmrp in-policy`** command.

Variable	Value
""	Indicates the policy used in the current session. If you do not enter a name, you delete the policy. **Important:** If you delete an accept policy from an interface, VLAN, or port, you change the configuration; you do not delete the policy itself.
policy name	Specifies the name of the announce policy you create.
port	Identifies the number of the port.

Example of applying a DVMRP accept policy to a port

Procedure steps

Apply the policy named policy2 to port 1/5.

```
ERS_8606:5# config ether 1/5 ip dvmrp in-policy policy2
```

Applying the advertisement of local networks policy to an interface

Apply the advertisement of local networks policy to advertise local networks over an interface

Procedure steps

Advertise local networks:

```
config ip dvmrp interface <ipaddr> advertise-self enable
```

Variable definitions

Use the data in the following table to use the **`config ip dvmrp interface advertise-self`** command.

Variable	Value
enable	Enables the advertisement of local routes for the selected interface to other switches in the network.
disable	Disables the advertisement of local routes for the selected interface to other switches in the network.
ipaddr	Specifies the address of the interface.

Example of applying the advertisement of local networks policy to an interface

Procedure steps

Enable advertisement of local routes on interface 100.100.100.2

```
ERS_8606:5# config ip dvmrp interface 100.100.100.2 advertise-self enable
```

Applying the advertisement of local networks policy over a VLAN

Apply the advertisement of local networks policy to advertise local networks over a VLAN.

Procedure steps

Advertise local networks:

```
config vlan <vid> ip dvmrp advertise-self enable
```

Variable definitions

Use the data in the following table to use the **config vlan ip dvmrp advertise-self** command.

Variable	Value
enable	Enables the advertisement of local routes over the selected VLAN to other switches in the network.
disable	Disables the advertisement of local routes over the selected VLAN to other switches in the network.
vid	Specifies the number of the VLAN.

Example of applying the advertisement of local networks policy over a VLAN

Procedure steps

Enable the advertisement of local networks on VLAN 100.

```
ERS_8606:5# config vlan 100 ip dvmrp advertise-self enable
```

Applying the advertisement of local networks policy over a port

Apply the advertisement of local networks policy to advertise local networks over a port.

Procedure steps

Advertise local networks:

```
config ether <port> ip dvmrp advertise-self enable
```

Variable definitions

Use the data in the following table to use the **`config ether ip dvmrp advertise-self`** command.

Variable	Value
enable	Enables the advertisement of local routes over the selected port to other switches in the network.
disable	Disables the advertisement of local routes over the selected port to other switches in the network.
port	Specifies the number of the port.

Example of applying the advertisement of local networks policy over a port

Procedure steps

Enable the advertisement of local routes on port 9 of the card in slot 1.

```
ERS_8606:5# config ethernet 1/9 ip dvmrp advertise-self enable
```

Creating a passive interface

Create a passive interface to drop all types of incoming DVMRP packets from neighbors and not send out probes or route reports to neighbor switches.

Procedure steps

Create a passive DVMRP interface:

```
config ip dvmrp interface <ipaddr> create passive
```

Variable definitions

Use the data in the following table to use the **config ip dvmrp interface create** command.

Variable	Value
active	Specifies that the interface receives all types of incoming DVMRP packets from neighbors and sends out probes or route reports to its neighbor switches.
ipaddr	Specifies the address of the interface.
passive	Specifies that the interface drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to its neighbor switches.

Example of creating a passive interface

Procedure steps

Create a passive interface named 100.100.100.2.

```
ERS_8606:5# config ip dvmrp interface 100.100.100.2 create passive
```

Configuring an active or passive interface type

Configure an interface as active or passive.

Procedure steps

Configure an interface type:

```
config ip dvmrp interface <ipaddr> interface-type <active|
passive>
```

Variable definitions

Use the data in the following table to use the **config ip dvmrp interface interface-type** command.

Variable	Value
active	Specifies that the interface receives all types of incoming DVMRP packets from neighbors and sends out probes or route reports to its neighbor switches.
ipaddr	Specifies the address of the selected interface.
passive	Specifies that the interface drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to its neighbor switches.

Example of configuring an active or passive interface type

Procedure steps

Set interface 100.100.100.2 as active.

```
ERS_8606:5# config ip dvmrp interface 100.100.100.2 interface-type active
```

Configuring an active or passive VLAN type

Configure a VLAN interface as active or passive. A passive VLAN drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to neighbor switches.

Procedure steps

Configure a VLAN interface:

```
config vlan <vid> ip dvmrp interface-type <active|passive>
```

Variable definitions

Use the data in the following table to use the **config vlan ip dvmrp interface-type** command.

Variable	Value
active	Specifies that the interface receives all types of incoming DVMRP packets from neighbors and sends out probes or route reports to its neighbor switches.
passive	Specifies that the interface drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to its neighbor switches.
vid	Specifies a VLAN ID from 1–4092

Example of configuring an active or passive VLAN type

Procedure steps

Set VLAN 100 as passive.

```
ERS_8606:5# config vlan 100 ip dvmrp interface-type active
```

Configuring an active or passive port type

Configure a port interface as active or passive. A passive port drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to neighbor switches.

Procedure steps

Configure a port type:

```
config ethernet <port> ip dvmrp interface-type <active|passive>
```

Variable definitions

Use the data in the following table to use the **config ethernet ip dvmrp interface-type** command.

Variable	Value
active	Specifies that the interface receives all types of incoming DVMRP packets from neighbors and sends out probes or route reports to its neighbor switches.
passive	Specifies that the interface drops all types of incoming DVMRP packets from neighbors and does not send out probes or route reports to its neighbor switches.
port	Specifies the number of the port.

Example of configuring an active or passive port type

Procedure steps

Set port 9 in slot 1 as active.

```
ERS_8606:5# config ethernet 1/9 ip dvmrp interface-type active
```

Chapter 10: DVMRP configuration using the ACLI

Configure the Distance Vector Multicast Routing Protocol (DVMRP) is to exchange multicast routing information between routers.

For more information about DVRMP configuration examples, see *Technical Configuration Guide for Passport 8600 Multicast DVMRP*. You can find this document at http://www.avaya.com/support.

Prerequisites to DVMRP configuration

- Enable DVMRP globally.
- Disable Protocol Independent Multicast-Sparse Mode (PIM-SM) from the interface where you want to configure DVMRP because you cannot configure PIM-SM and DVMRP on the same interface.

  Important:

  Avaya recommends that you do not change the configuration from PIM to DVMRP, or from DVMRP to PIM, while multicast traffic is flowing on the network.
- A switch can use a mix of DVMRP and PIM-SM interfaces if it is configured as a multicast border router (MBR).

Important:

For DVMRP scaled configurations with more than one thousand streams, to avoid multicast traffic loss, increase routing protocol timeouts, for example, dead interval for OSPF.

DVMRP configuration navigation

- Job aid on page 232
- Showing DVMRP next hops on page 234
- Applying the default route policy to a VLAN on page 235
- Applying the default route policy to a port on page 237

- Creating a DVMRP policy on page 239
- Applying a DVMRP announce policy to a VLAN on page 246
- Applying a DVMRP announce policy to a port on page 247
- Applying a DVMRP accept policy to a VLAN on page 247
- Applying a DVMRP accept policy to a port on page 248
- Applying the advertisement of local networks policy over a VLAN on page 249
- Applying the advertisement of local networks policy over a port on page 250
- Configuring an active or passive VLAN type on page 251
- Creating an active port on page 251
- Configuring an active or passive port type on page 252

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	
show ip dvmrp next-hop	
Global Configuration mode	
ip dvmrp show-next-hop-table	
Route map Configuration mode	
route-map <policy-name> <seq_number>	**enable**
	match as-path Word<0-256>
	match community Word<0-256>
	match community-exact <enable| disable>
	match extcommunity Word<0-1027>
	match interface Word<0-259>
	match local-preference <0-2147483647>
	match metric Word<0-65535>

Command	Parameter
	match network Word<0-259>
	match next-hop Word<0-259>
	match protocol <Any|xxx>
	match route-source Word<0-259>
	match route-type <any|local|internal|external|external-1|external-2>
	match tag Word<0-256>
	vrf
	vrfids
	name <policy_name>
	<permit|deny>
	set as-path Word<0-256>
	set as-path-mode <tag|prepend>
	set automatic-tag <enable|disable>
	set community Word<0-256>
	set community-mode <additive|none|unchanged>
	set injectlist Word<0-1027>
	set ip-preference <0-255>
	set local-preference <0-65535>
	set mask <A.B.C.D>
	set metric <0-65535>
	set metric-type <type1|type2>
	set metric-type-internal <0|1>
	set next-hop <A.B.C.D>
	set nssa-pbit enable
	set origin <igp|egp|incomplete>
	set origin-egp-as <o0-65535>
	set tag <0-65535>

Command	Parameter
	set weight <0-65535>
Interface Configuration mode	
ip dvmrp	**active**
	advertise-self
	default-listen
	default-supply
	default-supply-metric <cost>
	enable
	in-policy <policy_name>
	interface-type active
	metric <cost>
	out-policy <policy_name>

Showing DVMRP next hops

Showing the next-hop table is disabled by default. Disable this setting to avoid using the large amount of memory required for these tables in a scaled multicast environment with a large number of VLANs.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable showing the next-hop table:

   ```
   ip dvmrp show-next-hop-table
   ```

2. Display information about the DVMRP next hops:

   ```
   show ip dvmrp next-hop
   ```

Job aid

The following table shows the field descriptions for this command.

Field	Description
SOURCE	Indicates the network address that, when combined with the corresponding value of dvmrpRouteNextHopSourceMask, identifies the sources for a next hop on an outgoing interface.
MASK	Indicates the network mask that, when combined with the corresponding value of dvmrpRouteNextHopSource, identifies the sources for a next hop on an outgoing interface.
INTERFACE	Indicates the outgoing interface for this next hop.
TYPE	Displays leaf if no downstream dependent neighbors exist on the outgoing virtual interface. Otherwise, the type is branch.

Applying the default route policy to a VLAN

Apply the default route policy to a VLAN. With DVMRP routing policies, you can improve the management of the DVMRP routing tables by controlling how the routing table is populated and how Avaya Ethernet Routing Switch 8800/8600 devices exchange the routes.

Prerequisites

- You must log on to the VLAN Interface Configuration mode.

Procedure steps

1. Apply the policy to an interface:

   ```
   ip dvmrp default-listen default-supply default-supply-metric
   <value>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
active	Enables DVMRP on a specific interface of a specific interface type. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is enabled.
advertise-self	Advertises the local network if the setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to advertise.
default-listen	Learns the default route over the specified interface if this feature is enabled on the interface. The default setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to listen.
default-supply	Generates and advertises the default route if this feature is enabled on the interface. The default setting is disable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to not supply.
default-supply-metric <*cost*>	Advertises the specified metric over the interface if you configured the interface to supply the default route. The range is 1–31 hops. The default setting is 1 hop. To set this option to the default value, use the **default** operator with the command.
enable	Enables DVMRP on a specific interface. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
in-policy <*policy_name*>	Specifies the DVMRP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Variable	Value
interface-type active	Specifies the interface type as active. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is active.
metric <*value*>	Specifies the DVMRP route metric value. • *cost* is the metric value with a range of 1–31. To set this option to the default value, use the **default** operator with the command. The default is 1.
out-policy <*policy_name*>	Specifies the DVMRP route out-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Applying the default route policy to a port

Apply the default route policy to a port. With DVMRP routing policies, you can improve the management of the DVMRP routing tables by controlling how the routing table is populated and how Avaya Ethernet Routing Switch 8800/8600 devices exchange the routes.

Prerequisites

- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

1. Apply the policy to an interface:

   ```
   ip dvmrp default-listen default-supply default-supply-metric
   <cost>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
active	Enables DVMRP on a specific interface of a specific interface type. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is enabled.
advertise-self	Advertises the local network if the setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to advertise.
default-listen	Learns the default route over the specified interface if this feature is enabled on the interface. The default setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to listen.
default-supply	Generates and advertises the default route if this feature is enabled on the interface. The default setting is disable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is to supply.
default-supply-metric <*cost*>	Advertises the specified metric over the interface if you configured the interface to supply the default route. The range is 1–31 hops. The default setting is 1 hop. To set this option to the default value, use the **default** operator with the command.
enable	Enables DVMRP on a specific interface. Use the **no** operator to remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
in-policy <*policy_name*>	Specifies the DVMRP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Variable	Value
interface-type active	Specifies the interface type as active. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is active.
metric <*cost*>	Specifies the DVMRP route metric. • *cost* is the metric value with a range of 1–31. To set this option to the default value, use the **default** operator with the command. The default is 1.
out-policy <*policy_name*>	Specifies the DVMRP route out-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Creating a DVMRP policy

Before you can apply an announce or accept policy to an interface, VLAN, or port, you must first create and configure the policy.

Prerequisites

- You must configure an IP prefix list to use in the policy.
- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Create a DVMRP policy:

   ```
   route-map <policy-name> <seq_number>
   ```

   You are now in Route-map Configuration mode.
2. Create an action for the policy:

   ```
   route-map <policy-name> <permit|deny> <seq_number>
   ```

3. Create a match condition by using one or more of the match parameters listed in the following variable definitions table with the following command:

   ```
   route-map <policy-name> <seq_number>
   ```

4. Enable the policy:

   ```
   route-map <policy-name> seq <seq_number> enable
   ```

5. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **route-map** command.

Variable	Value
enable	Enables a route policy with a policy name and a sequence number. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
match as-path *Word<0–256>*	If configured, the switch matches the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified as-lists. This field is used only for BGP routes and ignored for all other route types. • *as-list* specifies the list IDs of up to four as-lists, separated by commas. Use the **no** operator to later remove this configuration.
match community *Word<0–256>*	If configured, the switch matches the community attribute of the BGP routes against the contents of the specified community-lists. This field is used only for BGP routes and ignored for all other route types. • *community-list* specifies the list IDs of up to four defined community-lists, separated by commas. Use the **no** operator to later remove this configuration.
match community-exact <enable|disable>	When disabled, match-community results in a match when the community attribute of the BGP routes matches an entry of a community-list specified in match-community. When enabled, match-community results in a match when the community attribute of the BGP routes

Variable	Value
	matches all of the entries of all the community-lists specified in match-community. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
match extcommunity Word<0–1027>	Matches the community in the community-id list where **word** is between 0–1027. Represent multiple community-id list as 2,4,5,6,7.
match interface *Word<0–259>*	If configured, the switch matches the IP address of the interface from which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other route types. • *prefix-list* specifies the names of up to four defined prefix lists, separated by commas. Use the **no** operator to later remove this configuration.
match local-preference <0-2147483647>	Matches the preference value from 0–2147483647. To set this option to the default value, use the **default** operator with the command. The default is 0.
match metric *<0-65535>*	If configured, the switch matches the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored. • *metric* is 0–65535. The default is 0. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command.
match network *Word<0–259>*	If configured, the switch matches the destination network against the contents of the specified prefix lists. • *prefix-list* specifies the names of up to four defined prefix lists, separated by commas. Use the **no** operator to later remove this configuration.

Variable	Value
match next-hop *Word<0–259>*	If configured, matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes. • *prefix-list* specifies the names of up to four defined prefix lists, separated by commas. Use the **no** operator to later remove this configuration.
match protocol *<Any|xxx>*	If configured, matches route policy to the protocol from which the route is learned. xxx is local, ospf, ebgp, ibgp, rip, dvmrp, static, or a combination separated by a vertical bar (|). This field is used only for Routing Information Protocol (RIP) announcement purposes. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is any.
match route-source *Word<0–259>*	If configured, matches the next-hop IP address for RIP, BGP, and DVMRP routes and advertising router IDs for Open Shortest Path First (OSPF) routes against the contents of the specified prefix list. This option is ignored for all other route types. • *prefix-list* specifies the names of up to four defined prefix lists, separated by commas. Use the **no** operator to later remove this configuration.
match route-type *<any|local|internal|external|external-1|external-2>*	Configures a specific route-type to match (applies only to OSPF routes). Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is any.
match tag *Word<0–256>*	Specifies a list of tags that is used during the match criteria process. Contains one or more tag values. • *tag* is a value from 0–256. Use the **no** operator to later remove this configuration.
vrf	Match the VRF name.
vrfids	Match a range of VRFs.
name *<policy_name>*	Renames a policy after its creation. This command changes the name field for all sequence numbers under the policy.

Variable	Value
<permit|deny>	Specifies the action to take when a policy matches a specific route. Permit accepts the route and deny ignores the route. To set this option to the default value, use the **default** operator with the command. The default is permit.
policy-name	Indicates the name of the specified policy, which is a string length from 1–64 characters.
seq_number	Indicates the number of the specified policy, which is a number from 1–65535.
set as-path *Word<0–256>*	If configured, the switch adds the as-number of the as-list to the BGP routes that match this policy. • *as-list-id* specifies the list ID of up to four defined as-lists, separated by commas. Use the **no** operator to later remove this configuration.
set as-path-mode <tag|prepend>	**prepend** is the default configuration. The switch prepends the as-number of the as-list specified in set-as-path to the old as-path attribute of the BGP routes that match this policy. Use the **no** operator to later remove this configuration.
set automatic-tag <enable| disable>	Configures the tag automatically. This option is used for BGP routes only. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is disabled.
set community *Word<0–256>*	If configured, the switch adds the community number of the community list to the BGP routes that match this policy. • *community-list* specifies the list ID of up to four defined community lists, separated by commas. Use the **no** operator to later remove this configuration.
set community-mode <additive| none|unchanged>	Configures the community mode. • **additive**—the switch prepends the community number of the community list specified in set-community to the old community path attribute of the BGP routes that match this policy. • **none**—the switch removes the community path attribute of the BGP routes that match this policy. • **unchanged**—keeps the community attribute in the route path.

Variable	Value
	Use the **no** operator to later remove this configuration. The default is unchanged.
set injectlist *Word<0–1027>*	If configured, the switch replaces the destination network of the route that matches this policy with contents of the specified prefix list. • *prefix-list* specifies one prefix list by name. Use the **no** operator to later remove this configuration.
set local-preference *<0-65535>*	A value used during a route decision process in the BGP protocol. Applicable to BGP only. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is 0.
set mask *<A.B.C.D>*	If configured, the switch configures the mask of the route that matches this policy. This applies only to RIP accept policies. • *ipaddr* is a valid contiguous IP mask. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is 0.0.0.0.
set metric *<0-65535>*	If configured, the switch configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or default-import-metric is used. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command.
set metric-type *<type1|type2>*	If configured, sets the metric type for the routes that match this policy to announce into the OSPF domain. The default is type 2. This variable is applicable only for OSPF announce policies. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command.
set metric-type-internal	Sets the internal metric type, 0 or 1. To set this option to the default value, use the **default** operator with the command. The default is 0.
set next-hop *<A.B.C.D>*	Specifies the IP address of the next-hop router. Ignored for DVMRP routes. Use the **no** operator to later remove this configuration. To set this option to the

Variable	Value
	default value, use the **default** operator with the command. The default is 0.0.0.0.
set nssa-pbit enable	Configures the not-so-stubby-area (NSSA) translation P bit. This variable is applicable only for OSPF announce policies. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is enable.
set origin <*igp*|*egp*|*incomplete*>	If configured, the switch changes the origin path attribute of the BGP routes that match this policy to the specified value. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is unchanged (not configured).
set origin-egp-as <*0-65535*>	Indicates the remote autonomous system number. This variable is applicable to BGP only. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is 0.
set tag <*0-65535*>	Configures the tag of the destination routing protocol. If you do not specify a tag, the switch forwards the tag value in the source routing protocol. A value of zero indicates that this parameter is not set. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is 0.
set weight <*0-65535*>	The weight value for the routing table. For BGP, this value overrides the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. This parameter is applicable to BGP only. A value of zero indicates that this parameter is not set. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is 0.

Example of creating a DVMRP announce policy

Procedure steps

1. Create the policy.

```
ERS-8606:5(config)# route-map policy2 1
```

2. Configure the action.

```
ERS-8606:5(config)# route-map policy2 deny 1
```

3. Set a match condition.

```
ERS-8606:5(config)# route-map policy2 1 match network prefix1
```

4. Enable the policy.

```
ERS-8606:5(config)# route-map policy2 1 enable
```

Applying a DVMRP announce policy to a VLAN

Apply a DVMRP announce policy to a VLAN to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Prerequisites

- You must create the policy before you can apply it to an interface.
- You must log on to the VLAN Interface Configuration mode.

Procedure steps

Apply a DVMP announce policy:

```
ip dvmrp out-policy <policy_name>
```

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
out-policy *<policy_name>*	Specifies the DVMRP route out-policy. • *policy_name* is a policy name that uses a string length of 0 to 64. To set this option to the default value, use the **default** operator with the command.

Applying a DVMRP announce policy to a port

Apply a DVMRP announce policy to a port to control which routes the switch sends to neighboring routers, to reduce the size of routing tables, or to provide a level of security for the network.

Prerequisites

- You must create the policy before you can apply it to an interface.
- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

Apply a DVMRP announce policy:

```
ip dvmrp out-policy <policy-name>
```

Variable definitions

Use the data in the following table to use the **`ip dvmrp`** command.

Variable	Value
out-policy <*policy_name*>	Specifies the DVMRP route out-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **`default`** operator with the command.

Applying a DVMRP accept policy to a VLAN

Apply a DVMRP accept policy to a VLAN to control the way DVMRP manages incoming route advertisements.

Prerequisites

- You must create the policy before you can apply it to an interface.
- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

Apply an accept policy:

```
ip dvmrp in-policy <policy-name>
```

Variable definitions

Use the data in the following table to use the **`ip dvmrp`** command.

Variable	Value
in-policy <*policy_name*>	Specifies the DVMRP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **`default`** operator with the command.

Applying a DVMRP accept policy to a port

Apply a DVMRP accept policy to a port to control the way DVMRP manages incoming route advertisements.

Prerequisites

- You must create the policy before you can apply it to an interface.
- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

Apply an accept policy:

```
ip dvmrp in-policy <policy-name>
```

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
in-policy <*policy_name*>	Specifies the DVMRP route in-policy. • *policy_name* is a policy name that uses a string length of 0–64. To set this option to the default value, use the **default** operator with the command.

Applying the advertisement of local networks policy over a VLAN

Apply the advertisement of local networks policy to advertise local networks over a VLAN.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

Advertise local networks:

```
ip dvmrp advertise-self
```

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
advertise-self	Advertises the local network if the setting is enable. Use the **no** operator to later remove this configuration. To set this option to the default value, use the **default** operator with the command. The default is enabled.

Applying the advertisement of local networks policy over a port

Apply the advertisement of local networks policy to advertise local networks over a port.

Prerequisites

- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

Advertise local networks:

```
ip dvmrp advertise-self
```

Variable definitions

Use the data in the following table to use the **ip dvmrp** command.

Variable	Value
advertise-self	Advertises the local network if the setting is enabled. Use the **no** operator to later remove this configuration. To set this option

Variable	Value
	to the default value, use the **default** operator with the command. The default is enabled.

Configuring an active or passive VLAN type

Configure a VLAN interface as active or passive.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Configure an interface type as active following:

   ```
   ip dvmrp interface-type active
   ```

2. Configure an interface type as passive:

   ```
   no ip dvmrp interface-type active
   ```

Creating an active port

Create an active port to receive all types of incoming DVMRP packets from neighbors, and send out probes or route reports to neighbor switches.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI.

Procedure steps

Create an active interface:

```
ip dvmrp active
```

Example of creating an active interface

Procedure steps

Create an active interface.

```
ERS-8606:5(config-if)# ip dvmrp active
```

Configuring an active or passive port type

Configure a port interface as active or passive.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI.

Procedure steps

1. Configure an interface type as active:

   ```
   ip dvmrp interface-type active
   ```

2. Configure an interface type as passive:

   ```
   no ip dvmrp interface-type active
   ```

Chapter 11: PIM configuration using Enterprise Device Manager

The Avaya Ethernet Routing Switch 8800/8600 supports two modes of Protocol Independent Multicast (PIM): Sparse Mode (SM) and Source Specific Multicast (SSM).

- PIM-SM supports multicast groups spread out across large areas of a company or the Internet.
- PIM-SSM optimizes PIM-SM by simplifying the many-to-many model (servers-to-receivers).

For more information about PIM-SM concepts and terminology, see IP multicast fundamentals on page 21.

Prerequisites to PIM configuration

- Before you can configure PIM-SM, you must configure an IP interface. For more information about IP interfaces, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing,* (NN46205-523).
- Disable the Distance Vector Multicast Routing Protocol (DVMRP) on the interface where you want to configure PIM-SM because you cannot configure PIM-SM and DVMRP on the same interface.

  Important:

  Avaya recommends that you do not change the configuration from PIM to DVMRP, or from DVMRP to PIM, while multicast traffic is flowing on the network.

  A switch can use a mix of DVMRP and PIM-SM interfaces if it is configured as an multicast border router (MBR).
- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM.

  For more information about RIP and OSPF, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — OSPF and RIP, (*NN46205-522).

  PIM-SM requires a unicast protocol to use to multicast traffic within the network when it performs the Reverse Path Forwarding (RFP) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that PIM-enabled routers use to communicate. The unicast routing table must contain a

route to every multicast source in the network as well as routes to PIM entities like the rendezvous point (RP) and bootstrap router (BSR).

- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- Configure one or more RPs for the groups that are used by a multicast application in the network.

Important:

You must configure and enable PIM on the circuitless IP interface before you configure PIM on the RP. For more information about how to configure a PIM-SM RP for a circuitless IP interface, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — BGP Services,* (NN46205-510).

- Configure one or more BSRs to propagate RP information to all switches in the network.
- If you connect a PIM-SM domain to a DVMRP domain, configure the switch connecting the domains as a multicast border router (MBR) switch with the corresponding PIM-SM interfaces enabled with PIM-SM, and the DVMRP interfaces enabled with DVMRP.

Important:

Routes to sources in a PIM domain must not use a lower cost through the DVMRP domain to ensure that multicast routing from these sources works properly. Configure MBR switches with this design guideline in mind.

Navigation

- Enabling static RP on page 255
- Configuring static RP on page 255
- Configuring a candidate bootstrap router on page 257
- Viewing the current BSR information on page 258
- Changing the VLAN interface type on page 260
- Editing PIM interface parameters on page 260
- Viewing PIM-SM neighbor parameters on page 262
- Viewing the RP set parameters on page 263
- Configuring a candidate RP on page 264
- Enabling square-SMLT globally on page 265

Enabling static RP

Enable static RP to avoid the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism.

Prerequisites

- To configure PIM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.
3. Click **PIM**.
4. Click **Mode: sm** (sparse mode).
5. Click **Enable**.
6. Click **Apply**.

   Information messages appear to remind you that traffic does not stop immediately, and that RP information learned through the BSR is lost.

   Important:

   Because you cannot configure a static RP-enabled switch as a BSR, the Current BSR tab disappears from this dialog box after you click Apply.
7. Click **Yes** to continue.

Configuring static RP

With static RP, you can configure a static entry for an RP. When configured, static RP ignores the BSR mechanism and uses the statically configured RPs only. A static RP-enabled switch uses this feature to communicate with switches from other vendors that do not use the BSR mechanism.

Prerequisites

- Before you can configure a static RP, you must enable the following:
  - PIM-SM
  - static RP
- To configure PIM on a specific VRF instance, first change the VRF instance as required.
- After meeting the prerequisites, keep in mind the following configuration considerations:
  - You cannot configure a static RP-enabled switch as a BSR or as a candidate RP (C-RP router).
  - All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.
  - Static RPs do not age. They cannot time out.
  - Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless it is configured with that static RP.
  - Configure all the switches in the network (including switches from other vendors) to map to the same RP.
  - In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interfaces as an RP.
  - To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of Avaya and other vendor switches across the network, ensure that all switches or routers use the same active RP because other vendors use different algorithms to elect the active RP. The Avaya Ethernet Routing Switch 8800/8600 uses the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.
  - Static RP configured on the switch is active as long as the switch uses a unicast route to the network for the static RP. If the switch loses this route, the static RP is invalidated, and the hash algorithm is invoked to remap all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm is invoked to remap the affected groups.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Static RP** tab.

4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Variable definitions

Use the data in the following table to configure the PIM Static RP tab.

Variable	Value
GroupAddress	The IP address of the multicast group. When combined with the group mask, it identifies the range of the multicast addresses that the RP handles.
GroupMask	The address mask of the multicast group. When combined with the group address, it identifies the range of the multicast addresses that the RP handles.
Address	The IP address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid when the switch uses a unicast route to the network for the static RP and is invalid otherwise.

Configuring a candidate bootstrap router

PIM-SM cannot operate without a BSR. Although a PIM-SM domain can use only one active BSR, you can configure additional routers as candidate BSRs (C-BSR). C-BSRs provide backup protection in case the primary BSR fails.

The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **PIM** tab.

5. Click **Enable**.
6. In the **CBSRPreference** box, type the preference.

   The C-BSR with the highest BSR-preference and address becomes the active BSR. The default is –1, which indicates that the current interface is not a C-BSR.
7. Click **Apply**.

Viewing the current BSR information

View the current BSR information to review the configuration.

Prerequisites

- To view and configure the Current BSR tab, you must disable StaticRP in the IP, PIM, and Globals tab. After you disable static RP, it results in the loss of the static RP configuration information.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Current BSR** tab.

Variable definitions

Use the data in the following table to use the PIM Current BSR tab.

Variable	Value
Address	Shows the IP address of the current BSR for the local PIM domain.
FragmentTag	Shows a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same fragment tag.
HashMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. With the hashmask, a small number of consecutive groups can always hash to the same RP.

Variable	Value
Priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
BootStrapTimer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

Configuring a PIM virtual neighbor

A virtual neighbor is a PIM neighbor IP address on the Avaya Ethernet Routing Switch 8800/8600 neighbor table. Use a virtual neighbor where the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Use the following procedure to configure a PIM virtual neighbor.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Virtual Neighbor** tab.
4. Click **Insert**.
5. Specify the PIM virtual neighbor address and the associated local interface.
6. Click **Insert**.

Variable definitions

Use the data in the following table to configure a PIM virtual interface.

Variable	Value
Address	Specifies the address of the PIM virtual neighbor.
Interface	Specifies the local interface associated with the PIM virtual neighbor.

Changing the VLAN interface type

Change the state (active or passive) of PIM on a VLAN interface.

Prerequisites

- Before you change the state of PIM on a VLAN interface, you must first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when streams are received.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select the VLAN ID that you want to configure with PIM.
4. Click **IP** from the menu bar.
5. Click the **PIM** tab.
6. Ensure that **Enable** is not selected.

   Tip:

   If a check mark appears next to **Enable**, clear the check box to disable PIM.
7. Click **Apply**.
8. Reenable PIM on the VLAN interface.

Editing PIM interface parameters

Edit PIM parameters for an interface to customize your PIM configuration.

Prerequisites

- Before you change the state (active or passive) of a PIM interface using the Interface Type field, first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when streams are received.
- To configure PIM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Interfaces** tab.
4. Edit the fields by clicking on them, and then select the new values.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the PIM Interfaces tab.

Variable	Value
IfIndex	Shows the interface Index.
Address	Shows the IP address of the PIM interface.
NetMask	Shows the network mask for the IP address of the PIM interface.
Mode	Shows the configured mode of this interface. The valid modes are SSM and sparse. This field is a read-only field.
InterfaceType	Specifies if the interface is active or passive.
DR	Shows the router with the highest IP address on a LAN designated to perform these tasks.
HelloInterval	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default is 30 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default is 60 seconds.

Variable	Value
CBSRPreference	Configures your preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OperState	Indicates the status of PIM on this interface: enabled or disabled.

Viewing PIM-SM neighbor parameters

View PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Prerequisites

- To view PIM information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Neighbors** tab.

Variable definitions

Use the data in the following table to use the PIM Neighbors tab.

Variable	Value
Address	Shows the IP address of the PIM neighbor.
IfIndex	Shows the slot or port number or VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Shows the elapsed time since this PIM neighbor last became a neighbor of the local router.
ExpiryTime	Shows the time remaining before this PIM neighbor times out.

Viewing the RP set parameters

The RP set is a list of rendezvous point addresses. The BSR constructs this list from C-RP advertisements and then distributes it to all PIM routers in the PIM domain for the BSR. View the parameters for troubleshooting purposes.

Prerequisites

- To view PIM information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **RP Set** tab.

Variable definitions

Use the data in the following table to use the PIM RP Set tab.

Variable	Value
Component	Displays a number that uniquely identifies the component. Each protocol instance that connects to a separate domain uses a different index value.
GroupAddress	Shows the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GroupMask	Shows the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
Address	Shows the IP address of the C-RP router.
HoldTime	Shows the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.

Variable	Value
ExpiryTime	Shows the time remaining before this C-RP router times out.

Configuring a candidate RP

Configure a C-RP router to add it to the RP Set.

You can configure only one interface on an Avaya Ethernet Routing Switch 8800/8600 for multiple groups; that is, you cannot configure multiple interfaces for multiple groups.

Using the GroupMask value, you can configure a candidate RP for several groups in one configuration. For example, with a C-RP configuration with a GroupAddress value of 224.0.0.0 and a GroupMask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Prerequisites

- To configure PIM-SM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Candidate RP** tab.
4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Variable definitions

Use the data in the following table to configure the PIM Candidate RP tab.

Variable	Value
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
InterfaceAddress	Configures the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

Important:

The following procedure also enables full-mesh configurations.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Multicast Square SMLT** tab.
4. To enable square-SMLT, select the **MulticastSquareSmltEnable** box.
5. Click **Apply**.

Chapter 12: PIM configuration using the CLI

This section describes the commands used to configure Protocol Independent Multicast (PIM) on your Avaya Ethernet Routing Switch 8800/8600. PIM provides two modes: Sparse Mode (SM) and Source Specific Multicast (SSM).

For more information about PIM-SM configuration examples, see *Technical Configuration Guide for PP8600 PIM-SM*. You can find this document at http://www.avaya.com/support.

Prerequisites to PIM configuration

- Configure an IP interface. For more information about IP interfaces, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing,* (NN46205-523).
- Disable the Distance Vector Multicast Routing Protocol (DVMRP) from the interface where you want to configure PIM-SM because you cannot configure PIM-SM and DVMRP on the same interface.

  Important:

  Avaya recommends that you do not change the configuration from PIM to DVMRP, or from DVMRP to PIM, while multicast traffic is flowing on the network.

  - A switch can use a mix of DVMRP and PIM-SM interfaces if it is configured as an multicast border router (MBR).
- Configure a unicast protocol, for example Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) globally and on the interfaces where you want to configure PIM-SM. For more information about RIP and OSPF, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — OSPF and RIP,* (NN46205-522).
- PIM-SM on an Avaya Ethernet Routing Switch 8800/8600 requires the following configuration:
  - Enable PIM-SM globally.
  - Enable PIM-SM on individual interfaces.
  - Configure one or several RPs for the groups that are used by a multicast application in the network.

    Important:

    You must first configure and enable PIM on the circuitless IP interface before you configure PIM on the rendezvous point (RP). For more information about how to

configure a PIM-SM RP for a circuitless IP interface, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — BGP Services,* (NN46205-510).

- Configure one or several bootstrap routers to propagate RP information to all switches in the network.
- If you connect a PIM-SM domain to a DVMRP domain, configure the switch interconnecting the domains as an MBR switch with the corresponding PIM-SM interfaces enabled with PIM-SM, and the DVMRP interfaces enabled with DVMRP.

Important:

Routes to sources in a PIM domain must not use a lower cost through the DVMRP domain to ensure that multicast routing from these sources works properly. Configure MBR switches with this design guideline in mind.

PIM configuration navigation

- Job aid on page 268
- Enabling a PIM multicast border router on page 270
- Configuring the PIM interface virtual neighbor on page 271
- Configuring a candidate rendezvous point on page 272
- Configuring static RP on page 274
- Configuring a candidate BSR on an interface on page 276
- Configuring a candidate BSR on an Ethernet port on page 277
- Configuring a candidate BSR on a VLAN on page 278
- Enabling square-SMLT globally on page 279

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ethernet <ports> ip pim candbsr`	`info`
	`disable`

Command	Parameter
	enable preference <value>
config ip pim candbsr interface <ipaddr>	info
	enable preference <value>
	disable
config ip pim candrp	info
	add grp <value> mask <value> rp <value>
	delete grp <value> mask <value>
config ip pim interface <ipaddr> virtual-neighbor	info
	add <ipaddr>
	delete <ipaddr>
config ip pim mbr	info
	disable
	enable
config ip pim static-rp	info
	add grp <value> mask <value>rp <value>
	delete grp <value> mask <value> rp <value>
	disable
	enable
	specific-route <enable | disable>
config ip vrf <vrfName> pim candbsr interface <ipaddr>	info
	enable preference <value>
	disable
config ip vrf <vrfName> pim candrp	info
	add grp <value> mask <value> rp <value>
	delete grp <value> mask <value>

Command	Parameter
`config ip vrf <vrfName> pim interface <ipaddr> virtual-neighbor`	`add <ipaddr>`
	`delete <ipaddr>`
	`info`
`config ip vrf <vrfName> pim static-rp`	`info`
	`add grp <value> mask <value>rp <value>`
	`delete grp <value> mask <value> rp <value>`
	`disable`
	`enable`
	`specific-route <enable | disable>`
`config sys mcast-smlt`	`info`
	`square-smlt <enable|disable>`
`config vlan <vid> ip pim candbsr`	`info`
	`disable`
	`enable preference <value>`
`show ip pim candidate-rp`	
`show ip pim rp-set`	

Enabling a PIM multicast border router

PIM MBRs connect PIM domains to other multicast routing domains and to the rest of the Internet. Enable the MBR functionality on an Avaya Ethernet Routing Switch 8800/8600 to connect a PIM-SM domain to a DVMRP domain.

Procedure steps

Enable a PIM MBR:

```
config ip pim mbr enable
```

Variable definitions

Use the data in the following table to use the **config ip pim mbr** command.

Variable	Value
disable	Disables PIM MBR on the switch.
enable	Enables PIM MBR on the switch.
info	Displays the current PIM MBR configuration setting.

Configuring the PIM interface virtual neighbor

A virtual neighbor is a PIM neighbor IP address on the Avaya Ethernet Routing Switch 8800/8600 neighbor table. Use a virtual neighbor where the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

You configure PIM on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

Configure the PIM interface virtual neighbor:

```
config ip pim interface <ipaddr> virtual-neighbor add <ipaddr>
```

Variable definitions

Use the data in the following table to use the **config ip pim interface virtual-neighbor** and **config ip vrf <vrfName> pim interface virtual-neighbor** commands.

Variable	Value
add <*ipaddr*>	Adds the virtual neighbor on the local switch interface. • *ipaddr* is the IP address of the selected interface.

Variable	Value
delete <*ipaddr*>	Deletes the virtual neighbor on the local switch interface. • *ipaddr* is the IP address of the selected interface.
info	Displays current virtual neighbor settings on the local switch interface.
ipaddr	Indicates the IP address of the selected interface.

Configuring a candidate rendezvous point

Configure a candidate rendezvous point (C-RP) to serve as backup to the RP router.

You can configure only one interface on an Avaya Ethernet Routing Switch 8800/8600 for multiple groups. You cannot configure multiple interfaces for multiple groups.

Avaya recommends that you configure a candidate rendezvous point with a multicast group address and mask as close to the multicast source as possible. You can also use the mask value to configure a C-RP router for several groups in one configuration. For example, a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0 permits you to configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

You configure PIM on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Procedure steps

1. Configure a candidate rendezvous point:

   ```
   config ip pim candrp add grp <value> mask <value> rp <value>
   ```

2. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ip pim candidate-rp
   ```

Variable definitions

Use the data in the following table to use the **`config ip pim candrp`** and **`config ip vrf <vrfName> pim candrp`** commands.

Variable	Value
add grp <*value*> mask <*value*> rp <*value*>	Adds a candidate RP to the RP Set. • **`add grp`** <*value*> is the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router. • **`mask`** <*value*> is the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router. • **`rp`** <*value*> is the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.
delete grp <*value*> mask <*value*>	Deletes a candidate RP from the RP set. • **`delete grp`** <*value*> is the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router. • **`mask`** <*value*> is the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
info	Displays current RP configuration settings on the local router interface.

Job aid

The following table shows the field descriptions for this command.

Field	Description
GRPADDR	The IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	The address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RPADDR	The IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Configuring static RP

With static RP, you can configure a static entry for an RP. When configured, static RP ignores the BSR mechanism and uses the statically configured RPs only. Static RP-enabled switches use this feature to communicate with switches from other vendors that do not use the BSR mechanism.

Important:

You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.

All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

You configure PIM on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Prerequisites

- Enable PIM-SM globally.

Procedure steps

1. Enable static RP:

   ```
   config ip pim static-rp enable
   ```

   The following message appears:

   ```
   WARNING: RP information learnt dynamically through BSR
   functionality will be lost.

   Do you wish to enable Static RP? (y/n) ?
   ```

2. Configure static RP:

   ```
   config ip pim static-rp add grp <value> mask <value> rp
   <value>
   ```

3. Configure all the switches in the network (including switches from other vendors) to map to the same RP.

Variable definitions

Use the data in the following table to use the **`config ip pim static-rp`** and **`config ip vrf <vrfName> pim static-rp`** command.

Variable	Value
add grp <*value*> mask <*value*> rp <*value*>	Adds a static RP entry to the RP set. • **`grp`** <*value*> is the IP address of the multicast group. When combined with the group mask, it identifies the range of the multicast addresses that the RP handles. • **`mask`** <*value*> is the address mask of the multicast group. When combined with the group address, it identifies the range of the multicast addresses that the RP handles. • **`rp`** <*value*> is the IP address of the static RP.
delete grp <*value*> mask <*value*> rp <*value*>	Deletes a static RP entry from the RP set. • **`grp`** <*value*> is the IP address of the multicast group. When combined with the group mask, it identifies the range of the multicast addresses that the RP handles. • **`mask`** <*value*> is the address mask of the multicast group. When combined with the group address, it identifies the range of the multicast addresses that the RP handles. • **`rp`** <*value*> is the IP address of the static RP.
disable	Disables static RP on the switch.
enable	Enables static RP on the switch.
info	Displays current PIM settings on the switch.
specific-route <enable | disable>	With static RP enabled, if the route to the RP is removed, the Avaya Ethernet Routing Switch 8800/8600 can fail over to an alternate static RP. However, if a default route exists in the routing table, that default route still appears as an active route to the failed RP. In this case, the switch does not fail over to the alternate RP. A similar situation exists with SMLT-based configurations, where an internal-only default static route is used during IST failover and recovery. In this case, the internal default route appears as an active

Variable	Value
	route to the failed RP, and therefore does not failover to the alternate RP. To resolve these situations, you can configure the lookup for static RP to be chosen from the specific route rather than the best route. In this case, when the route to the active RP fails, the switch no longer interprets the default route as a valid route for RP purposes, and therefore fails over to the alternate RP.

Example of configuring static RP

Procedure steps

1. Enable static RP.

```
ERS-8606:5#config ip pim static-rp
ERS-8606:5/config/ip/pim/static-rp# enable
```

```
WARNING: RP information learnt dynamically through BSR
functionality will be lost.

Do you wish to enable Static RP? (y/n) ? y
```

2. Add a static RP entry.

```
ERS-8606:5/config/ip/pim/static-rp# add grp 239.255.0.0 mask 255.255.0.0
rp 100.1.1.1
```

Configuring a candidate BSR on an interface

PIM-SM cannot run without a BSR. Although a PIM-SM domain can use only one active BSR, you can configure additional routers as candidate BSRs (C-BSR). C-BSRs provide backup protection in the event that the primary BSR fails.

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

You configure PIM on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Procedure steps

Configure a candidate BSR:

```
config ip pim candbsr interface <ipaddr> enable preference
<value>
```

Variable definitions

Use the data in the following table to use the **`config ip pim candbsr interface`** and **`config ip vrf <vrfName> pim candbsr interface`** command.

Variable	Value
disable	Disables the C-BSR on this interface.
enable preference <*value*>	Enables the C-BSR on this interface and configures its preference value to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
info	Displays the C-BSR preference setting for this interface.
ipaddr	Indicates the IP address of the selected interface.

Configuring a candidate BSR on an Ethernet port

PIM-SM cannot run without a BSR. Although a PIM-SM domain can use only one active BSR, you can configure additional routers as candidate BSRs (C-BSR). C-BSRs provide backup protection in the event that the primary BSR fails.

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure steps

Configure a candidate BSR:

```
config ethernet <ports> ip pim candbsr enable preference
<value>
```

Variable definitions

Use the data in the following table to use the **`config ethernet ip pim candbsr`** command.

Variable	Value
disable	Disables the C-BSR on this interface.
enable preference <*value*>	Enables the C-BSR on this interface and configures its preference value to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
info	Displays the C-BSR preference setting for this interface.
ports	Specifies the port using the convention {slot/port[-slot/port][,...]}.

Configuring a candidate BSR on a VLAN

PIM-SM cannot run without a BSR. Although a PIM-SM domain can use only one active BSR, you can configure additional routers as candidate BSRs (C-BSR). C-BSRs provide backup protection in case the primary BSR fails.

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure steps

Configure a candidate BSR on a VLAN:

```
config vlan <vid> ip pim candbsr enable preference <value>
```

Variable definitions

Use the data in the following table to use the **config vlan ip pim candbsr** command.

Variable	Value
disable	Disables the C-BSR on this interface.
enable preference <*value*>	Enables the C-BSR on this interface and configures its preference value to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
info	Displays the C-BSR preference setting for this interface.
vid	Specifies a VLAN ID from 1–4092.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

Important:

The following command also enables full-mesh configurations.

Procedure steps

Enable square-SMLT:

```
config sys mcast-smlt square-smlt <enable|disable>
```

Variable definitions

Use the data in the following table to use the **config sys mcast-smlt** command.

Variable	Value
info	Displays the current square-SMLT configuration.

Variable	Value
square-smlt <enable|disable>	Enables or disables square-SMLT configuration. The default is disable.

Chapter 13: PIM configuration using the ACLI

This section describes the commands used to configure Protocol Independent Multicast (PIM) on your Avaya Ethernet Routing Switch 8800/8600. PIM provides two modes: Sparse Mode (SM) and Source Specific Multicast (SSM).

For more information about PIM-SM configuration examples, see *Technical Configuration Guide for PP8600 PIM-SM*. You can find this document at http://www.avaya.com/.

Prerequisites to PIM configuration

- Configure an IP interface. For more information about IP interfaces, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing,* (NN46205-523).
- Disable the Distance Vector Multicast Routing Protocol (DVMRP) on the interface where you want to configure PIM-SM, because you cannot configure PIM-SM and DVMRP on the same interface.

  Important:

  Avaya recommends that you do not change the configuration from PIM to DVMRP, or from DVMRP to PIM, while multicast traffic is flowing on the network.

  - A switch can use a mix of DVMRP and PIM-SM interfaces if it is configured as an multicast border router (MBR).
- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM. For more information about RIP and OSPF, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — OSPF and RIP,*(NN46205-522).
- To configure PIM-SM on an Avaya Ethernet Routing Switch 8800/8600, you require the following configuration:
  - Enable PIM-SM globally.
  - Enable PIM-SM on individual interfaces.
  - Configure one or several RPs for the groups that are used by a multicast application in the network.

Important:

You must first configure and enable PIM on the circuitless IP interface before configuring PIM on the rendezvous point (RP). For more information about how to configure PIM-SM RP for a circuitless IP interface, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — BGP Services,* (NN46205-510).

- Configure one or more BSRs to propagate RP information to all switches in the network.
- If you connect a PIM-SM domain to a DVMRP domain, configure the switch interconnecting the domains as an MBR switch with the corresponding PIM-SM interfaces enabled with PIM-SM, and the DVMRP interfaces enabled with DVMRP.

Important:

Routes to sources in a PIM domain must not use a lower cost through the DVMRP domain to ensure that multicast routing from these sources works properly. Configure MBR switches with this design guideline in mind.

PIM configuration navigation

- Job aid on page 282
- Enabling a PIM multicast border router on page 283
- Configuring the PIM virtual neighbor on page 284
- Configuring a candidate rendezvous point on page 285
- Configuring static RP on page 286
- Configuring a candidate BSR on an Ethernet port on page 288
- Configuring a candidate BSR on a VLAN on page 289
- Enabling square-SMLT globally on page 290

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	

Command	Parameter
show ip pim rp-candidate [vrf Word<0-16>] [vrfids Word<0-255>]	
Global Configuration mode	
ip pim	**mbr**
	rp-candidate group <A.B.C.D> <mask> rp <A.B.C.D>
	static-rp <A.B.C.D/X> <A.B.C.D>
	virtual-neighbor <ipaddr> <ipaddr>
multicast smlt-square	
Interface Configuration mode	
ip pim	**bsr-candidate preference <value>**
VRF Router Configuration mode	
ip pim	**rp-candidate group <A.B.C.D> <mask> rp <A.B.C.D>**
	static-rp <A.B.C.D/X> <A.B.C.D>
	virtual-neighbor <ipaddr> <ipaddr>

Enabling a PIM multicast border router

PIM MBRs connect PIM domains to other multicast routing domains and to the rest of the Internet. Enable the MBR functionality on an Avaya Ethernet Routing Switch 8800/8600 to connect a PIM-SM domain to a DVMRP domain.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Enable a PIM MBR:

```
ip pim mbr
```

Configuring the PIM virtual neighbor

A virtual neighbor is a PIM neighbor IP address on the Avaya Ethernet Routing Switch 8800/8600 neighbor table. You typically use a virtual neighbor when the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

You configure PIM on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Configure the PIM virtual neighbor:

```
ip pim virtual-neighbor <ipaddr> <ipaddr>
```

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
ipaddr	The first IP address indicates the IP address of the selected interface.
ipaddr	The second IP address Indicates the IP address of the neighbor.

Configuring a candidate rendezvous point

Configure a candidate rendezvous point (C-RP) to serve as backup to the RP router.

You can configure only one interface on an Avaya Ethernet Routing Switch 8800/8600 for multiple groups. You cannot configure multiple interfaces for multiple groups.

With the mask value, you can configure a C-RP router for several groups in one configuration. For example, a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0 permits you to configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

You configure PIM on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Add a candidate rendezvous point:

   ```
   ip pim rp-candidate group <A.B.C.D> <mask> rp <A.B.C.D>
   ```

2. Remove a candidate rendezvous point:

   ```
   no ip pim rp-candidate group <A.B.C.D> <mask>
   ```

3. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ip pim rp-candidate [vrf Word<0-16>] [vrfids Word<0-
   255>]
   ```

Variable definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
group <A.B.C.D>	Specifies the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
mask	Specifies the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
rp <A.B.C.D>	Specifies the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Job aid

The following table shows the field descriptions for this command.

Field	Description
GRPADDR	Displays the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Displays the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RPADDR	Displays the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Configuring static RP

With static RP, you can configure a static entry for an RP. When configured, static RP ignores the bootstrap router (BSR) mechanism and uses the statically configured RPs only. Static RP-enabled switches use this feature to communicate with switches from other vendors that do not use the BSR mechanism.

You configure PIM on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Important:

You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.

All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Prerequisites

- Enable PIM-SM globally.
- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable static RP:

   ```
   ip pim static-rp
   ```

2. Enter `y` at the following prompt:

   ```
   WARNING: RP information learnt dynamically through BSR
   functionality will be lost. Do you wish to enable Static RP?
   (y/n) ?
   ```

3. Configure a static RP entry:

   ```
   ip pim static-rp <A.B.C.D/X> <A.B.C.D>
   ```

4. Configure all the switches in the network (including switches from other vendors) to map to the same RP.

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
A.B.C.D/X	Specifies the IP address and address mask of the multicast group. When combined, the IP address and address mask identify the range of the multicast addresses that the RP handles.
A.B.C.D	Specifies the IP address of the static RP.

Example of configuring static RP

Procedure steps

Add a static RP entry.

```
ERS-8606:5(config)# ip pim static-rp 239.255.0.0/255.255.0.0 100.1.1.1
```

Configuring a candidate BSR on an Ethernet port

PIM-SM cannot run without a BSR. Although a PIM-SM domain can use only one active BSR, you can configure additional routers as candidate BSRs (C-BSR). C-BSRs provide backup protection in the event that the primary BSR fails.

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Prerequisites

- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI.

Procedure steps

Configure a candidate BSR:

```
ip pim bsr-candidate preference <value>
```

Variable definitions

Use the data in the following table to use the **`ip pim`** command.

Variable	Value
preference <*value*>	Enables the C-BSR on this interface and configures its preference value, from 1–255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates

Variable	Value
	that the current interface is not a C-BSR. To set this option to the default value, use the **default** operator with the command.

Configuring a candidate BSR on a VLAN

PIM-SM cannot run without a BSR. Although a PIM-SM domain can use only one active BSR, you can configure additional routers as candidate BSRs (C-BSR). C-BSRs provide backup protection in case the primary BSR fails.

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

Configure a candidate BSR on a VLAN:

```
ip pim bsr-candidate preference <value>
```

Variable definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
preference <*value*>	Enables the C-BSR on this interface and configures its preference value, from 1–255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. To set this option to the default value, use the **default** operator with the command.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

Important:

The following command also enables full-mesh configurations.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Enable square-SMLT:

```
multicast smlt-square
```

Chapter 14: IGMP configuration using Enterprise Device Manager

Hosts use Internet Group Management Protocol (IGMP) to report their multicast group memberships to neighboring multicast routers.

For more information about troubleshooting, see *Avaya Ethernet Routing Switch 8800/8600 Troubleshooting,* (NN46205-703).

Prerequisites to IGMP configuration

- Configure IGMP on a Layer 3 interface by first enabling multicast routing, for example, Distance Vector Multicast Routing Protocol (DVMRP), Protocol Independent Multicast-Sparse Mode (PIM-SM), or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

Navigation

- Enabling IGMP snoop on a VLAN on page 292
- Configuring IGMP interface static members on page 292
- Configuring the SSM channel table on page 294
- Configuring the SSM range and global parameters on page 295
- Configuring multicast stream limitation on an interface on page 297
- Configuring multicast stream limitation on a VLAN on page 298
- Configuring multicast stream limitation on an Ethernet port on page 299
- Configuring multicast stream limitation members on page 300
- Adding a multicast stream limitation member on page 301
- Deleting a multicast stream limitation member on page 302
- Editing the IGMP interface table on page 302
- Configuring IGMP sender entries on page 305

- Configuring fast leave mode on page 306
- Configuring multicast access control for an interface on page 308
- Viewing IGMP cache information on page 309
- Viewing multicast router discovery information on page 310
- Viewing IGMP snoop information on page 312
- Viewing IGMP group information on page 313

Enabling IGMP snoop on a VLAN

Enable IGMP snopping on a VLAN to optimize the multicast data flow for a group within a VLAN to only members of the group that uses IGMP snoop. The switch listens to group reports from each port and builds a database of multicast group members for each port. It suppresses the reports heard by not forwarding them to other hosts, forcing the members to continuously send their own reports. The switch relays group membership from the hosts to the multicast routers. It also forwards queries from multicast routers to all port members of the VLAN. Furthermore, it multicasts data only to the participating group members and to the multicast routers within the VLAN.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP** from the menu bar.
5. Click **IGMP**.
6. Click **SnoopEnable**.
7. Click **ProxySnoopEnable**.
8. In the **StreamLimitEnable** checkbox, select **enable**.
9. In the **Maximum Number of Stream**, type the maximum number of streams.
10. Click **Apply**.

Configuring IGMP interface static members

Configure IGMP interface static members to add members to a snoop group. Some sources do not join a multicast group before transmitting a multicast stream. When this situation occurs

and if no other group members joined in the VLAN, the data is flooded to all port members of the VLAN. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams always forward to the multicast router within the VLAN, in addition to the ports you configure for this static entry.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Important:

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Static** tab.
4. In the **Static** tab, click **Insert**.
5. Edit the appropriate data.
6. Click **Insert**.

Variable definitions

Use the data in the following table to configure the Insert Static dialog box.

Variable	Value
IfIndex	Shows the interface where the IGMP entry is enabled.
GrpAddr	Configures the multicast group address of the multicast stream. Within the indicated valid range (224.0.1.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.

Variable	Value
MemberPorts	Specifies the ports to which you want to redirect the multicast stream for this multicast group. The ports must be member ports of the VLAN.
NotAllowedToJoin	Specifies the ports that do not receive the multicast stream for this multicast group.

Configuring the SSM channel table

The SSM channel table consists of entries that map groups to their sending source. SSM channels cannot conflict with static source groups and vice versa. After you configure an SSM channel or a static source group, the switch performs a consistency check to make sure no conflicts exist. You cannot map one group (G) to different sources for both a static source group and an SSM channel.

The consistency check applies to all SSM channel entries, even if they are disabled. If you disable an entry, it becomes inactive; you do not delete the entry and you can later reenable it.

After you disable an SSM channel, the Avaya Ethernet Routing Switch 8800/8600 stops multicast traffic from the specified source to the specified group. You can use this static setting as a security feature to block traffic from a certain source to a specific group.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **SsmChannel** tab.
4. Click **Insert**.
5. Type the IP address for the multicast group and source.
6. Click **Insert**.

   The SsmChannel tab appears with the entry you just created in the table.

Tip:

You can change the default status of an SSM channel from enable to disable by clicking in the AdminState field.

Variable definitions

Use the data in the following table to configure the SsmChannel tab.

Variable	Value
IpMulticastGrp	An IP multicast address that is within the SSM range.
IpSource	The IP address of the source that is sending traffic to the group.
LearningMode	Displays whether the entry is statically configured (Static) or dynamically-learned from IGMPv3 (Dynamic). This a read-only field.
Activity	Displays the current activity of the selected (S,G) entry. True indicates that traffic is flowing to the switch, otherwise, it appears as false. This a read-only field for the Avaya Ethernet Routing Switch 8800/8600.
AdminState	The administrative state for the selected static entry. This state determines whether the switch uses the static entries. Set this field to enable (default) to use the entry or disable to save for future use.

Configuring the SSM range and global parameters

Use the SSM range parameter to extend the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations.

The other global parameters in this dialog box enable the IGMPv3 dynamic learning feature and set the administrative state for all the entries in the SSM channel table.

Important:

If you change the RangeGroup configuration, the switch reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), this procedure also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

Prerequisites

- To change the RangeGroup configuration, you must first disable PIM.
- To change the RangeGroup configuration, you must delete all entries in the SSM channel table before you configure the new IP multicast group address.
- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **SsmGlobal** tab.
4. Configure the appropriate fields.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the SsmGlobal tab.

Variable	Value
DynamicLearning	Enables the dynamic learning of SSM channel (S,G) pairs from IGMPv3 reports. As new SSM channels are learned, they appear in the SSM channel table.
AdminAction	Configures the administrative state, which determines whether the switch uses the table entries: • none (default)—Does not set the administrative state globally so that you can set it for individual SSM channel table entries. • enableAll—Globally activates all the static entries in the SSM channel table. This setting does not affect the dynamically learned entries. • disableAll—Globally inactivates all the static entries in the SSM channel table. This setting does not affect the dynamically learned entries.

Variable	Value
RangeGroup	Configures the IP multicast group address. The lowest group address is 224.0.1.0 and the highest is 239.255.255.255. The default is 232.0.0.0.
RangeMask	Configures the address mask of the multicast group. The default is 255.0.0.0.

Configuring multicast stream limitation on an interface

Use multicast stream limitation to limit the number of concurrent multicast streams on the interface. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After the stream limit is met, the interface drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a set limit of multicast streams on an interface.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit** tab.
4. To change the status of an interface, click the **StreamLimitEnable** field for the interface, and then select one of the following from the menu:
   - **enable**
   - **disable**

   If the interface is enabled, you can edit the **Maximum Number of Stream** field.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the StreamLimit tab.

Variable	Value
Interface	Displays the slot or port number or VLAN ID for this interface.
Stream Limit Enable	Enables or disables stream limitation on this interface.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this interface. The range is from 0–65535, and the default is 4.
Current Number Of Stream	Displays the current number of streams received on this interface. This value is a read-only value.

Configuring multicast stream limitation on a VLAN

Use multicast stream limitation to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After the stream limit is met, the VLAN drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a set limit of multicast streams on an interface.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP** from the menu bar.
5. Select **IGMP**.
6. Configure the appropriate settings.
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the VLAN stream limitation settings.

Variable	Value
StreamLimitEnable	Enables or disables stream limitation on this VLAN.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this VLAN. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.

Configuring multicast stream limitation on an Ethernet port

Use multicast stream limitation to limit the number of concurrent multicast streams on the port. Limit the number of streams to protect the bandwidth on a specific interface and control access to multicast streams.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **IGMP** tab.
5. In the **StreamLimitEnable** field, click the **Enable** option button.
6. In **Maximum Number of Stream**, type the number of streams.
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Port stream limitation fields.

Variable	Value
StreamLimitEnable	Enables or disables stream limitation on this port.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this port. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.

Configuring multicast stream limitation members

Configure multicast stream limitation members on ports of the specified interface to set the maximum number of streams on the interface.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.
4. In the **MaxStreams** box, type the new number.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the StreamLimit Members tab.

Variable	Value
IfIndex	Displays the name of the VLAN.
Port	Lists each slot or port number for this interface with stream limitation enabled.
MaxStreams	Configures the maximum number of allowed streams for this specific port. The number of allowed streams cannot

Variable	Value
	exceed the maximum number for the interface. The range is from 0–65535 and the default is 4.
NumStreams	Displays the current number of streams received on this interface. This value is a read-only value.

Adding a multicast stream limitation member

Add a multicast stream limitation member to an interface.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.
4. Click **Insert**.
5. Do one of the following:
   - Type the number of the VLAN to which you want to add a member.
   - Select a VLAN from the **Vlan ID** list.
6. Do one of the following:
   - Type the number of the slot or port that you want to add as a member.
   - Click the ellipsis [**...**] button and select a slot or port from the graphic display.

   Important:

   You must select one of the ports in the VLAN that you selected in the previous step.
7. Do one of the following:
   - Type a maximum number of streams.

- Accept the default of 4.

8. Click **Insert**.

Deleting a multicast stream limitation member

Delete a multicast stream limitation member from an interface to remove it from the configuration.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.
4. Click the row that lists the member you want to delete.
5. Click **Delete**.

Editing the IGMP interface table

Use the Interface tab to view or edit the IGMP interface table. When an interface does not use an IP address, it does not appear in the IGMP table. When an interface uses an IP address, and you do not enable DVMRP or PIM-SM, notInService appears in the Status field.

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Interface** tab.
4. Edit the appropriate information.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Interface tab.

Variable	Value
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The range is from 1–65535 and the default is 125.
Status	Shows the IGMP row status. When an interface the status is notInService.
Version	Configures the version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	Shows the version of IGMP currently running on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface is attached.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. This value is not configurable for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255, and the default is 100 tenths of a second (equal to 10 seconds).

Variable	Value
	Important: You must configure this value lower than the QueryInterval.
WrongVersionQueries	Shows the number of queries received with an IGMP version that does not match the interface. You must configure all routers on a LAN to run the same version of IGMP. If queries are received with the wrong version, this value indicates a version mismatch.
Joins	Shows the number of times a group membership is added on this interface; that is, the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	Tunes for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If a network is expected to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that one query for each query interval is dropped without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of second. Avaya recommends that you configure this parameter to values greater than 3. If a fast leave process is not required, Avaya recommends values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
FlushAction	Configures the flush action to one of the following: • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set.

Variable	Value
	Important: To maximize your network performance, Avaya recommends that you set this parameter according to the version of IGMP currently in use. • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
SnoopQuerierEnable	Enables snoop querier.
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
MvrEnable	Enables or disables MCast-VLAN-registration on a VLAN. It also specifies a VLAN works as a MCast-VLAN. A maximum of 16 MCast-VLAN are supported.
MvrProxyEnable	Enables or disables MCast-VLAN-registration proxy on a VLAN. You must enable MCast-VLAN-registration first.

Configuring IGMP sender entries

Configure IGMP sender entries to identify a source that sends multicast data to a multicast group.

Prerequisites

- To configure PIM-SM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Sender** tab.

4. Change the appropriate options.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Sender tab.

Field	Description
GrpAddr	Specifies the multicast group address of the multicast stream. Within the indicated valid range (224.0.1.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.
IfIndex	Specifies the interface where the IGMP entry is enabled.
MemberAddr	Specifies the IP address of a host.
Action	Flushes an entry or a group.
TPort	Identifies the T port.
State	Indicates whether a sender exists because of an IGMP access filter. The options are filtered and not filtered.

Configuring fast leave mode

Configure fast leave mode to control all IGMP fast leave enabled interfaces. Fast leave relies on an alternative leave process where the switch stops sending traffic for the group immediately after receiving a leave message, without issuing a query to check if other group members exist on the network. Use this global parameter to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces. It does not apply to Internet Group membership Authentication Protocol (IGAP) interfaces.

Prerequisites

- To configure PIM-SM on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Global** tab.
4. In **FastLeaveMode** box, do one of the following:
   - Click **multipleUser** to remove the leaving member or to stop traffic if no more member are present.
   - Click **oneUser** to stop first leave traffic.
5. Click enable in the **GenerateTrap** box to enable DVMRP to generate trap.
6. Click enable in the **GenerateLog** box to enable the DVMRP to generate log.
7. Click enable in the **Mvr** box to enable MCast-VLAN-registration globally.
8. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Global tab.

Variable	Value
multipleUser	Removes from the group only the IGMP member who sent the leave message. Traffic is not stopped if other receivers exist on the interface port. This value is the default.
oneUser	Removes all group members on a fast leave enabled interface port upon receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.
FastLeaveMode	Sets the fast-leave mode for all fast-leave enabled IGMP interfaces.
GenerateTrap	Enables the DVMRP to generate trap.
GenerateLog	Enables the DVMRP to generate log.

Variable	Value
Mvr	Enables MCast-VLAN-registration globally.

Configuring multicast access control for an interface

Configure multicast access control for a selected IGMP interface or a VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Prerequisites

- To configure IGMP on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Access Control** tab.
4. Click **Insert**.
5. Do one of the following:
   - Type the number of the slot, port, or VLAN ID that you want to add as a member.
   - Click the [**...**] button and select one from the graphic display.
6. Select the interface or VLAN.
7. Click **Ok**.
8. Click the ellipsis button [**...**] next to **PrefixListId**.
9. Select a prefix list ID/name.
10. Type the host address and host mask.
11. Select the action mode that you want for the specified host.
12. Click **Insert**.
13. Click **Close**.

Variable definitions

Use the data in the following table to configure the Access Control tab.

Variable	Value
IfIndex	Specifies the interface where the IGMP entry is enabled.
PrefixListId	Specifies a numeric string that identifies the prefix list.
HostAddr	Specifies the IP address of the host.
HostMask	Specifies the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
PrefixListName	Shows the name of the prefix list.
ActionMode	Specifies the action for the host identified by HostAddr. The options include the following: • denied IP multicast transmitted traffic (deny-tx) • denied IP multicast received traffic (deny-rx) • denied both IP multicast transmitted and received traffic (deny-both) • allowed IP multicast transmitted traffic (allow-only-tx) • allowed IP multicast received traffic (allow-only-rx) • allowed both IP multicast transmitted and received traffic (allow-only-both)

Viewing IGMP cache information

View IGMP cache information to view the group for which members exist on a specific interface.

Prerequisites

- To view IGMP information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Cache** tab.

Variable definitions

Use the data in the following table to use the Cache tab.

Variable	Value
Address	Shows the IP multicast group address for this entry that contains this information.
IfIndex	Shows the interface from which the corresponding multicast group address is heard.
LastReporter	Shows the IP address of the source of the last membership report received for this IP multicast group address on this interface. If no membership report is received, the object uses the value 0.0.0.0.
ExpiryTime	Shows the amount of time (in seconds) remaining before this entry is aged out.
Version1Host Timer	Shows the time remaining until the local router assumes that no IGMPv1 members exist on the IP subnet attached to the interface. Upon hearing IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local router ignores IGMPv2 leave messages for this group that it receives on this interface.

Viewing multicast router discovery information

View multicast router discovery information to view the current configuration.

Prerequisites

- To view IGMP information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Multicast Router Discovery** tab.

Variable definitions

Use the data in the following table to use the Multicast Router Discovery tab.

Variable	Value
Interface	Shows the interface where IGMP is enabled.
MrdiscEnable	Enables or disables the router interface to listen for multicast router discovery messages to determine where to send multicast source data and IGMPv2 reports. After you enable snoop, multicast router discovery is automatically enabled.
DiscoveredRouterPorts	Lists ports discovered by the IGMP Multicast Router Discovery (MRDISC) protocol. **Important:** The MRDISC protocol is not supported on brouter ports.
MaxAdvertiseInterval	Shows the maximum time allowed between sending router advertisements from the interface, in seconds. The range is from 2–180 seconds. The default is 20 seconds.
MinAdvertiseInterval	Shows the minimum time allowed between sending unsolicited router advertisements from the interface, in seconds. This value must be more than 3 seconds but no greater than the value assigned to the MaxAdvertiseInterval value.

Variable	Value
MaxInitialAdvertiseInterval	Configures the maximum number (in seconds) of multicast advertisement intervals that you can configure on the switch.
MaxInitialAdvertisements	Configures the maximum number of initial multicast advertisements that you can configure on the switch.
NeighborDeadInterval	Shows the time interval (in seconds) before the router interface drops traffic after a user leaves the multicast group.

Viewing IGMP snoop information

View information about IGMP snoop to see the current configuration.

Prerequisites

- To view IGMP information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Snoop** tab.

Variable definitions

Use the data in the following table to use the Snoop tab.

Variable	Value
IInterface	Shows the VLAN ID for the VLAN.
SnoopEnable	Shows the status of IGMP snoop. IGMP snoop works only when a multicast router exists in the VLAN.
SsmSnoopEnable	Shows the status of SSM snoop.

Variable	Value
ProxySnoopEnable	Indicates the status of the IGMP report proxy feature. After you enable this feature, hosts forward reports to the multicast router once for each group for each query interval, or after new group information is available. After you disable this feature, all reports from different hosts are forwarded to multicast routers, and more than one group report can be forwarded for the same multicast group for each query interval. The default is enabled.
FastLeaveEnable	Shows the status of fast leave for this port.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports are attached to a multicast router, so the multicast data and group reports are forwarded to the router. **Important:** Configure this field only when you use multiple multicast routers that are not attached to one another, but are attached to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and this field is configured, a multicast loop forms.
SnoopActiveMRouterPorts	Shows the active multicast router ports. Active multicast router ports are ports directly attached to a multicast router. These ports include the querier port and all ports in the forwarding state that you configured as well as those that were dynamically learned through receiving queries.
SnoopMRouterExpiration	Indicates the time remaining before the multicast router ages out. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The query maximum response interval (obtained from the queries received) is used as the timer resolution.

Viewing IGMP group information

View information about IGMP groups to see the current group operation on the switch.

Prerequisites

- To view IGMP information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Groups** tab.

Variable definitions

Use the data in the following table to use the Groups tab.

Variable	Value
IpAddress	Multicast group address (Class D) that members can join. A group address can be the same for many incoming ports.
Members	IP address of a member that issues a group report for this group.
InPort	Unique value to identify a brouter interface or a logical interface (VLAN) that receives group reports from various members.
IfIndex	Unique value that identifies a physical interface or a logical interface (VLAN), which receives group reports from various sources.
Expiration	Time left before the group report expires on this port. This variable is updated upon receiving a group report.

Chapter 15: IGMP configuration using the CLI

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.

For more information about IGMP configuration examples, see *PP8600 Technical Configuration Guide for IGMP Access Control*. You can find this document at http://www.avaya.com/support.

Prerequisites to IGMP configuration

- Complete one of the following tasks:
  - Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
  - Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Distance Vector Multicast Routing Protocol (DVMRP), Protocol Independent Multicast-Sparse Mode (PIM-SM), or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

Important:

To drop IGMPv2 control packets that do not use the router alert option, use the **`config ip igmp interface`** command and enable the router-alert parameter.

Navigation

- Job aid on page 316
- Configuring multicast stream limitation on an Ethernet port on page 323
- Configuring multicast stream limitation on an interface on page 325
- Configuring interface multicast stream limitation members on page 327
- Configuring multicast stream limitation on a VLAN on page 328
- Configuring VLAN multicast stream limitation members on page 329
- Configuring IGMP multicast router discovery options on page 331

- Configuring IGMP multicast router discovery on a VLAN on page 332
- Configuring IGMP interface static members on page 333
- Configuring SSM dynamic learning and range group on page 336
- Changing the SSM range group on page 338
- Configuring the SSM channel table on page 339
- Configuring multicast access control for an IGMP interface on page 342
- Configuring multicast access control for an IGMP Ethernet port on page 343
- Configuring multicast access control for a VLAN on page 345
- Configuring fast leave mode on page 346
- Configuring IGMP fast leave members on a VLAN on page 347
- Enabling IGMP L2 Querier globally on page 348
- Configuring L2 snoop Querier address on page 349
- Configuring L2 Querier on an IP interface on page 350
- Configuring L2 querier message source on a interface on page 350
- Viewing IP IGMP interface on page 351
- Viewing IGMP snoop configuration result on page 352
- Viewing IGMP configuration result on page 352

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section

Important:

You must globally enable DVMRP or PIM multicasting on the switch for these commands to take effect.

Table 5: Job aid: Roadmap of IGMP CLI commands

Command	Parameter
`config ethernet <ports> ip igmp access-control <name>`	`info`
	`create <HostAddress> <HostMask> <deny-tx|deny-rx| deny-both|allow-only-tx|`

Command	Parameter
	allow-only-rx|allow-only-both>
	delete <HostAddress> <HostMask>
	mode <HostAddress> <HostMask> <deny-tx|deny-rx|deny-both| allow-only-tx|allow-only-rx| allow-only-both>
config ethernet <ports> ip igmp stream-limit	info
	enable
	disable
	max-streams <integer>
config ip igmp fast-leave-mode <one-user|multiple-user>	
config ip igmp interface <ipaddr> access-control <name>	info
	delete <HostAddress> <HostMask>
	create <HostAddress> <HostMask> <deny-tx|deny-rx| deny-both|allow-only-tx| allow-only-rx|allow-only-both>
	mode <HostAddress> <HostMask> <deny-tx|deny-rx|deny-both| allow-only-tx|allow-only-rx| allow-only-both>
config ip igmp interface <ipaddr> mrdisc	info
	max-advertisement-interval <seconds>
	max-initial advertisement-interval <seconds>
	max-initial-advertisements <integer>
	min-advertisement-interval <seconds>

Command	Parameter
	mrdisc-enable <enable| disable>
	neighbor-dead-interval <seconds>
config ip igmp interface <ipaddr> static-members <FromGroupAddress-ToGroupAddress>	info
	add <ports> <static|blocked>
	create <ports> <static| blocked>
	delete
	remove <ports> <static| blocked>
config ip igmp interface <ipaddr> stream-limit	info
	enable
	disable
	max-streams <integer>
config ip igmp interface <ipaddr> stream-limit-members	info
	enable <ports> max-streams <value>
	disable <ports>
	set <ports> max-streams <value>
config ip igmp ssm	info
	dynamic-learning <enable| disable>
	ssm-grp-range group <value> mask <value>
config ip igmp ssm ssm-channel	info
	create group <value> source <value>
	delete group <value>
	disable <all or group> [<GroupAddress>]

Command	Parameter
	enable <all or group> [<GroupAddress>]
config ip vrf <WORD 0-64> igmp fast-leave-mode <one-user| multiple-user>	
config ip vrf <WORD 0-64> igmp interface <ipaddr> access-control <name>	info
	delete <HostAddress> <HostMask>
	create <HostAddress> <HostMask> <deny-tx|deny-rx| deny-both|allow-only-tx| allow-only-rx|allow-only-both>
	mode <HostAddress> <HostMask> <deny-tx|deny-rx|deny-both| allow-only-tx|allow-only-rx| allow-only-both>
config ip vrf <WORD 0-64> igmp interface <ipaddr> mrdisc	info
	max-advertisement-interval <seconds>
	max-initial advertisement-interval <seconds>
	max-initial-advertisements <integer>
	min-advertisement-interval <seconds>
	mrdisc-enable <enable| disable>
	neighbor-dead-interval <seconds>
config ip vrf <WORD 0-64> igmp interface <ipaddr> static-members <FromGroupAddress-ToGroupAddress>	info
	add <ports> <static|blocked>
	create <ports> <static| blocked>
	delete

Command	Parameter
	remove <ports> <static| blocked>
config ip vrf <WORD 0-64> igmp interface <ipaddr> stream-limit	info
	enable
	disable
	max-streams <integer>
config ip vrf <WORD 0-64> igmp interface <ipaddr> stream-limit-members	info
	enable <ports> max-streams <value>
	disable <ports>
	set <ports> max-streams <value>
config ip vrf <WORD 0-64> igmp ssm	info
	dynamic-learning <enable| disable>
	ssm-grp-range group <value> mask <value>
config ip vrf <WORD 0-64> igmp ssm ssm-channel	info
	create group <value> source <value>
	delete group <value>
	disable <all or group> [<GroupAddress>]
	enable <all or group> [<GroupAddress>]
config vlan <vid> ip igmp access-control <name>	info
	create <HostAddress> <HostMask> {deny-tx|deny-rx| deny-both|allow-only-tx| allow-only-rx|allow-only-both>
	delete <HostAddress> <HostMask>
	mode <HostAddress> <HostMask> {deny-tx|deny-rx|deny-both|

Command	Parameter
	allow-only-tx|allow-only-rx| allow-only-both>
config vlan <vid> ip igmp fast-leave-members	info
	disable <ports>
	enable <ports>
config vlan <vid> ip igmp mrdisc	info
	min-advertisement-interval <seconds>
	max-initial-advertisement-interval <seconds>
	max-initial-advertisements <integer>
	max-advertisement-interval <seconds>
	mrdisc-enable <enable| disable>
	neighbor-dead-interval <seconds>
config vlan <vid> ip igmp static-members <FromGroupAddress-ToGroupAddress>	info
	add <ports> <static|blocked>
	create <ports> <static| blocked>
	delete
	remove <ports> <static| blocked>
config vlan <vid> ip igmp stream-limit	info
	enable
	disable
	max-streams <integer>
config vlan <vid> ip igmp stream-limit-members	info
	enable <ports> max-streams <value>
	disable <ports>

Command	Parameter
	set <ports> max-streams <value>
show ip igmp	access [vrf <WORD 0-64>] [vrfids <0-255>]
	cache [vrf <WORD 0-64>] [vrfids <0-255>]
	group [count] [memb-subnet <value>] [group <value>] [vrf <WORD 0-64>] [vrfids <0-255>]
	igap
	igap-counters [vlan <value>]
	info [vrf <WORD 0-64>] [vrfids <0-255>]
	interface [vrf <WORD 0-64>] [vrfids <0-255>]
	mrdisc [vrf <WORD 0-64>] [vrfids <0-255>]
	mrdisc-neighbors [vrf <WORD 0-64>] [vrfids <0-255>]
	router-alert [vrf <WORD 0-64>] [vrfids <0-255>]
	sender [count] [memb-subnet <value>] [group <value>] [vrf <WORD 0-64>] [vrfids <0-255>]
	show-all [file <value>] [vrf <WORD 0-64>] [vrfids <0-255>]
	snoop [vrf <WORD 0-64>] [vrfids <0-255>]
	snoop-trace [src <value>] [grp <value>] [vrf <WORD 0-64>] [vrfids <0-255>]
	ssm-channel [vrf <WORD 0-64>] [vrfids <0-255>]
	ssm-global [vrf <WORD 0-64>] [vrfids <0-255>]

Command	Parameter
	static [vrf <WORD 0-64>] [vrfids <0-255>]
	stream-limit-interface [vrf <WORD 0-64>] [vrfids <0-255>]
	stream-limit-port [vrf <WORD 0-64>] [vrfids <0-255>]
config vlan <vid> ip igmp snoop-querier	<enable|disable>
config vlan <vid> ip igmp snoop-querier-addr <ipaddr>	
config ip [vrf <name>] igmp interface <ipaddr> snoop-querier	<enable|disable>
config ip [vrf <name>] igmp interface <ipaddr> snoop-querier-addr <ipaddr>	
config ip [vrf <name>] igmp interface <ipaddr> info	
show ip igmp snoop [vrf <WORD 0-64>] [vrfids <0-255>	
show ip info igmp [<vid>][vrf <WORD 0-64>] [vrfids <0-255>	

Configuring multicast stream limitation on an Ethernet port

Configure multicast stream limitation on an Ethernet port to limit the number of concurrent multicast streams on the port. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the port drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Procedure steps

1. Enable multicast stream limitation:

   ```
   config ethernet <ports> ip igmp stream-limit enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ethernet ip igmp stream-limit`** command.

Variable	Value
enable	Enables stream limitation on this port.
disable	Disables stream limitation on this port.
info	Displays information about the stream limits set on this port.
max-streams <*integer*>	Configures the maximum number of allowed streams on this port. The range is from 0–65535 and the default is 4.
ports	Specifies the port using the convention {slot/port[-slot/port][, ...]}.

Example of configuring multicast stream limitation on an Ethernet port

Procedure steps

1. Enable multicast stream limitation on the Ethernet port 1/3.

   ```
   ERS-8606:5# config ethernet 1/3 ip igmp stream-limit enable
   ```

2. Set the maximum number of allowed streams to 8.

   ```
   ERS-8606:5/config/ethernet/1/3/ip/igmp/stream-limit# max-streams 8
   ```

Configuring multicast stream limitation on an interface

Configure multicast stream limitation on an interface to limit the number of concurrent multicast streams on the interface. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the interface drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Enable multicast stream limitation on an interface:

   ```
   config ip igmp interface <ipaddr> stream-limit enable
   ```

2. Configure the remaining parameters as required.
3. Display information about the interfaces where multicast stream limitation is enabled:

   ```
   show ip igmp stream-limit-interface [vrf <value>] [vrfids
   <value>]
   ```

4. Display multicast stream limitation information for the ports on a specific interface:

   ```
   show ip igmp stream-limit-port [vrf <value>] [vrfids <value>]
   ```

Variable definitions

Use the data in the following table to use the **config ip igmp interface stream-limit** and **config ip vrf <vrfName> igmp interface stream-limit** commands.

Variable	Value
enable	Enables stream limitation on this interface.
disable	Disables stream limitation on this interface.

Variable	Value
info	Displays information about the stream limits set on this interface.
ipaddr	Indicates the IP address of the selected interface.
max-streams <*integer*>	Configures the maximum number of allowed streams on this interface. The range is from 0–65535 and the default is 4.
[vrf <value>] [vrfids <value>]	Specifies the name of a VRF or a range of VRF IDs to include in the show command output. If you do not specify a VRF name or range of VRF IDs, the command results show for the Global Router.

Example of configuring multicast stream limitation on an interface

Procedure steps

1. Enable multicast stream limitation on the interface at IP address 10.0.6.2.

```
ERS-8606:5# config ip igmp interface 10.0.6.2 stream-limit enable
```

2. Set the maximum number of allowed streams to 8.

```
ERS-8606:5/config/ip/igmp/interface/10.0.6.2/stream-limit#max-streams 8
```

Job aid

The following table shows the field descriptions for the **`show ip igmp stream-limit-interface`** command used in this procedure.

Field	Description
INTERFACE	Indicates the interface IP address.
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

The following table shows the field descriptions for the **`show ip igmp stream-limit-port`** command used in this procedure.

Field	Description
INTERFACE	Indicates the interface IP address.
PORT	Indicates the port for the VLAN.

Field	Description
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

Configuring interface multicast stream limitation members

Configure multicast stream limitation members on ports of the specified interface to limit the number of multicast groups that can join a VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the interface drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Procedure steps

1. Configure multicast stream limitation members on an interface:

   ```
   config ip igmp interface <ipaddr> stream-limit-members enable
   <ports> max-streams <value>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ip igmp interface stream-limit-members`** and **`config ip vrf <vrfName> igmp interface stream-limit-members`** commands.

Variable	Value
enable <*ports*> max-streams <*value*>	Enables stream limitation and configures the maximum number of allowed streams for the specified ports on this interface. The number of allowed streams cannot exceed the maximum

Variable	Value
	number for the interface. The range is from 0–65535 and the default is 4.
disable <*ports*>	Disables stream limitation for the specified ports on this interface.
info	Displays information about the stream limit members set on individual ports for this interface.
ipaddr	Indicates the IP address of the selected interface.
set <*ports*> max-streams <*value*>	Configures the maximum number of allowed streams for the specified ports on this interface. The range is from 0–65535 and the default is 4.

Example of configuring interface multicast stream limitation members

Procedure steps

1. Enable multicast stream limitation on ports 1/3 to 1/8 and set the maximum allowed number of streams to 6 for these ports.

```
ERS-8606:5# config ip igmp interface 192.32.96.82 stream-limit-members
enable 1/3-1/8 max-streams 6
```

2. Set the maximum number of allowed streams for ports 1/3 to 1/8 to 8.

```
ERS-8606:5/config/ip/igmp/interface/192.32.96.82/stream-limit-members#
set 1/3-1/8 max-streams 8
```

Configuring multicast stream limitation on a VLAN

Configure multicast stream limitation members on a VLAN to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the VLAN drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Procedure steps

1. Enable multicast stream limitation on a VLAN:

   ```
   config vlan <vid> ip igmp stream-limit enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp stream-limit`** command.

Variable	Value
enable	Enables stream limitation on this VLAN.
disable	Disables stream limitation on this VLAN.
info	Displays information about the stream limits set on this VLAN.
max-streams <*integer*>	Configures the maximum number of allowed streams on this VLAN. The range is from 0–65535 and the default is 4.
vid	Specifies the VLAN ID from 1–4093.

Example of configuring multicast stream limitation on a VLAN

Procedure steps

1. Enable multicast stream limitation on VLAN 3.

   ```
   ERS-8606:5# config vlan 3 ip igmp stream-limit enable
   ```

2. Set the maximum number of allowed streams to 8.

   ```
   ERS-8606:5/config/vlan/3/ip/igmp/stream-limit# max-streams 8
   ```

Configuring VLAN multicast stream limitation members

Configure multicast stream limitation on ports of a VLAN to limit the number of multicast groups that can join a VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the VLAN drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Procedure steps

1. Configure multicast stream limitation members on a VLAN:

   ```
   config vlan <vid> ip igmp stream-limit-members enable <ports>
   max-streams <value>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp stream-limit-members`** command.

Variable	Value
enable *<ports>* max-streams *<value>*	Enables stream limitation and configures the maximum number of allowed streams for the specified ports on this VLAN. The number of allowed streams cannot exceed the maximum number for the VLAN. The range is from 0–65535, and the default is 4.
disable *<ports>*	Disables stream limitation for the specified ports on this VLAN.
info	Displays information about the stream limit members set on individual ports for this VLAN.
set *<ports>* max-streams *<value>*	Configures the maximum number of allowed streams for the specified ports on this VLAN. The range is from 0–65535 and the default is 4.
vid	Specifies a VLAN ID from 1–4093.

Example of configuring VLAN multicast stream limit members

Procedure steps

1. Enable multicast stream limitation on ports 1/3 to 1/8 and set the maximum allowed number of streams to 6 for this interface.

```
ERS-8606:5# config vlan 3 ip igmp/ stream-limit-members
enable 1/3-1/8 max-streams 6
```

2. Set the maximum number of allowed streams for ports 1/3 to 1/8 to 8.

```
ERS-8606:5/config/vlan/3/ip/igmp/stream-limit-members# set
1/3-1/8 max-streams 8
```

Configuring IGMP multicast router discovery options

Configure the multicast router discovery options to enable the automatic discovery of multicast-capable routers.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Enable multicast router discovery:

```
config ip igmp interface <ipaddr> mrdisc mrdisc-enable enable
```

2. Configure the remaining parameters as required.

Important:

The Multicast Router Discovery (MRDISC) protocol is not supported on brouter ports.

Variable definitions

Use the data in the following table to use the **config ip igmp interface mrdisc** and **config ip vrf <vrfName> igmp interface mrdisc** commands.

Variable	Value
info	Displays information about the multicast route discovery on the interface.
ipaddr	Indicates the IP address of the selected interface.
max-advertisement-interval <*seconds*>	Configures the maximum number (in seconds) between successive advertisements. The range is 2–180 and the default is 20.

Variable	Value
	For this change to take effect, you must save the configuration and reset the switch.
max-initial advertisement-interval <seconds>	Configures the maximum number (in seconds) between successive initial advertisements. The range is 2–180 and the default is 2. For this change to take effect, you must save the configuration and reset the switch
max-initial-advertisements <integer>	Configures the maximum number of initial multicast advertisements after initialization. The range is 2–15 and the default is 3. For this change to take effect, you must save the configuration and reset the switch
min-advertisement-interval <*seconds*>	Configures the minimum number (in seconds) between successive advertisements. The range is 3–180 and the default is 15. For this change to take effect, you must save the configuration and reset the switch
mrdisc-enable <enable|disable>	Enables or disables the multicast route discovery option. The default is disable.
neighbor-dead-interval <*seconds*>	Configures the multicast router discovery dead interval—the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down. • *seconds* is a value from 1–59. The default is 30.

Configuring IGMP multicast router discovery on a VLAN

Configure IGMP multicast router discovery on a VLAN to enable the automatic discovery of multicast-capable routers.

Procedure steps

1. Enable multicast route discovery:

   ```
   config vlan <vid> ip igmp mrdisc mrdisc-enable <enable|
   disable>
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp mrdisc`** command.

Variable	Value
info	Displays multicast router discovery parameters on the VLAN.
min-advertisement-interval <*seconds*>	Configures the minimum number (in seconds) between successive advertisements. The range is 3–180 and the default is 15. For this change to take effect, you must save the configuration and reset the switch.
max-initial-advertisement-interval <*seconds*>	Configures the maximum number (in seconds) between successive initial advertisements. The range is 2–180 and the default is 2. For this change to take effect, you must save the configuration and reset the switch.
max-initial-advertisements <*integer*>	Configures the maximum number of initial multicast advertisements after initialization. The range is 2–15 and the default is 3. For this change to take effect, you must save the configuration and reset the switch.
max-advertisement-interval <*seconds*>	Configures the maximum number (in seconds) between successive advertisements. The range is 2–180 and the default is 20. For this change to take effect, you must save the configuration and reset the switch.
mrdisc-enable <enable|disable>	Enables multicast router discovery on the VLAN. The default is disable.
neighbor-dead-interval <*seconds*>	Configures the multicast router discovery dead interval—the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down. • *seconds* is a value from 1–59. The default is 30.
vid	Specifies a VLAN ID from 1–4092.

Configuring IGMP interface static members

Configure IGMP interface static members to add members to a snoop group. Some sources do not join a multicast group before transmitting a multicast stream. When this situation occurs and if no other group members joined in the VLAN, the data is flooded to all port members of

the VLAN. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams always forward to the multicast router within the VLAN, in addition to the ports you configure for this static entry.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Procedure steps

1. Configure interface static members:

   ```
   config ip igmp interface <ipaddr> static-members
   <FromGroupAddress-ToGroupAddress>
   ```

2. Configure the parameters as required.

Variable definitions

Use the data in the following table to use the **`config ip igmp interface static-members`** and **`config ip vrf <vrfName> igmp interface static-members`** command.

Variable	Value
add <*ports*> <static|blocked>	Adds a static-member entry to the IGMP interface. • *ports* is the port or list of ports to which you want to redirect the multicast stream for this multicast group. • **`static|blocked`** configures the route to static or blocked.
create <*ports*> <static|blocked>	Creates static members on the interface. • *ports* is the port or list of ports to which you want to redirect the multicast stream for this multicast group. • *static|blocked* configures the route to static or blocked.
delete	Deletes the static members on the interface.
FromGroupAddress-ToGroupAddress	Indicates the IP address range {a.b.c.d[-w.x.y.z]} of the selected multicast group.

Variable	Value
info	Displays information about the static members of the VLAN.
ipaddr	Indicates the IP address of the selected interface.
remove <*ports*> <static|blocked>	Removes slots/ports from the static members of a protocol-based VLAN. • *ports* is the port or list of ports you want to remove from the multicast stream. • **`static|blocked`** configures the multicast entry to static or blocked.

Configuring IGMP static members on a VLAN

Configure IGMP static members to add members to a snoop group. Some sources do not join a multicast group before transmitting a multicast stream. When this situation occurs and if no other group members joined in the VLAN, the data is flooded to all port members of the VLAN. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. When you create the entry, multicast data streams always forward to the multicast router within the VLAN, in addition to the ports you configure for this static entry.

Procedure steps

1. Configure static members:

   ```
   config vlan <vid> ip igmp static-members <FromGroupAddress-
   ToGroupAddress>
   ```

2. Configure the parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp static-members`** command.

Variable	Value
add <*ports*> <static|blocked>	Adds a static-member entry to the VLAN.

Variable	Value
	• *ports* is the port or list of ports to which you want to redirect the multicast stream for this multicast group. • **`static|blocked`** configures the route to static or blocked.
create <*ports*> <static|blocked>	Creates a static member entry on the VLAN. • *ports* is the port or list of ports to which you want to redirect the multicast stream for this multicast group. • **`static|blocked`** configures the route to static or blocked.
delete	Deletes a static-member entry from the VLAN.
FromGroupAddress-ToGroupAddress	Indicates the IP address or range of the selected multicast group in the format {a.b.c.d[-w.x.y.z]}.
info	Displays information about the static members of the VLAN.
remove <*ports*> <static|blocked>	Removes a port from the static-member entry to the VLAN. • *ports* is the port or list of ports to which you want to redirect the multicast stream for this multicast group. • **`static|blocked`** configures the route to static or blocked.
vid	Specifies a VLAN ID from 1–4092.

Configuring SSM dynamic learning and range group

Configure SSM dynamic learning and a range group to enable the IGMPv3 dynamic learning feature and to extend the default SSM range of 232/8 to include an IP multicast address.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Prerequisites

- To define the range group, you must first disable PIM.

Procedure steps

1. Configure SSM dynamic learning:

   ```
   config ip igmp ssm dynamic-learning <enable|disable>
   ```

2. Configure the range group:

   ```
   config ip igmp ssm ssm-grp-range group <value> mask <value>
   ```

Variable definitions

Use the data in the following table to use the **`config ip igmp ssm`** and **`config ip vrf <vrfName> igmp ssm`** command.

Variable	Value
dynamic-learning <enable|disable>	Enables the dynamic learning of SSM channel (S,G) pairs from IGMPv3 reports. As new SSM channels are learned, they appear in the SSM channel table. The default is enable.
info	Displays the SSM range and the status of the SSM channel table entries.
ssm-grp-range group <*value*> mask <*value*>	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. • **`group`** <*value*> is an IP multicast address within the range of 224.0.1.0 and 239.255.255.255. The default is 232.0.0.0. • **`mask`** <*value*> is the IP address mask of the multicast group. The default is 255.0.0.0.

Example of configuring SSM dynamic learning and range group

Procedure steps

1. Disable PIM.

```
ERS-8606:5# config ip pim disable
```

2. Define the SSM range group address (234.0.0.0) and mask (255.0.0.0).

```
ERS-8606:5# config ip igmp ssm ssm-grp group 234.0.0.0 mask 255.0.0.0
```

```
WARNING: All Static Source Group entries in the SSM range
will be deleted

Do you wish to change SSM range (y/n) ? y
```

3. Enable dynamic learning from IGMPv3 reports.

```
ERS-8606:5# config ip igmp ssm dynamic-learning enable
```

4. Enable PIM.

```
ERS-8606:5# config ip pim enable
```

Changing the SSM range group

Change the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address.

Important:

This procedure reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), it also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap (BSR) is local.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Disable PIM:

```
config ip pim disable
```

If you forget to disable PIM, the following error message appears:

```
Error: PIM is enabled in SSM mode, disable PIM
```

2. Delete each entry in the SSM channel table:

   config ip igmp ssm-channel delete group <*GroupAddress*>

   If you forget to delete the SSM channels, the following error message appears:

```
Error: SSM source group table not empty
```

3. Configure the new IP multicast group address:

   config ip igmp ssm ssm-grp-range group <*value*> **mask** <*value*>

   The following message appears:

```
WARNING: All Static Source Group entries in the SSM range will be deleted
Do you wish to change SSM range (y/n) ?
```

4. Enter **Y**.
5. Enable PIM:

   `config ip pim enable`

Configuring the SSM channel table

Configure the SSM channel table to map groups to their sending source. SSM channels cannot conflict with static source groups and vice versa. After you configure an SSM channel or a static source group, the switch performs a consistency check to make sure no conflicts exist. You cannot map one group (G) to different sources for both a static source group and an SSM channel.

The consistency check applies to all SSM channel entries, even if they are disabled. If you disable an entry, it becomes inactive; you do not delete the entry and you can later reenable it.

After you disable an SSM channel, the Avaya Ethernet Routing Switch 8800/8600 stops multicast traffic from the specified source to the specified group. You can use this static setting as a security feature to block traffic from a certain source to a specific group.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

Configure the SSM channel table:

```
config ip igmp ssm ssm-channel
```

Variable definitions

Use the data in the following table to use the **`config ip igmp ssm ssm-channel`** and **`config ip vrf <vrfName> igmp ssm ssm-channel`** commands.

Variable	Value
create group <*value*> source <*value*>	Creates a static SSM channel table entry by specifying the group and source IP addresses. • **`group`** <*value*> is an IP multicast address within the SSM range defined by **`ssm-grp-range group`**. • **`source`** <*value*> is an IP host address that sends traffic to the group.
delete group <*value*>	Deletes the SSM channel table entry that you specify. • *value* is the IP multicast address of the table entry you want to delete.
disable <*all or group*> [<*GroupAddress*>]	Disables the administrative state for all static entries in the SSM channel table (all) or for a specific entry (group). This setting does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry. • *all* refers to all the static entries in the SSM channel table. • *group* requires the GroupAddress of the entry you want to disable.
enable <*all or group*> [<*GroupAddress*>]	Enables the administrative state for all static entries in the SSM channel table (all) or for a specific entry (group). This setting does not affect the dynamically learned entries.

Variable	Value
	This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry. • *all* refers to all the static entries in the SSM channel table. • *group* requires the GroupAddress of the entry you want to disable.
info	Displays the SSM range and the status of the SSM channel table entries.

Example of configuring the SSM channel table

Procedure steps

1. Create an SSM channel table entry for the multicast group 234.0.1.0 and the source at 192.32.99.151.

```
ERS-8606:5# config ip igmp ssm ssm-channel create group 234.0.1.0 source
192.32.99.151
```

2. Set the administrative state to enable all the static SSM channel table entries.

```
ERS-8606:5#config ip igmp ssm ssm-channel enable all
```

3. Use the **`info`** command to show a summary of the results.

```
ERS-8606:5# config ip igmp ssm ssm-channel info
```

Following is an example of the command output:

```
Sub-Context: clear config dump monitor show test trace wsm sam
Current Context:
create :
group : 234.0.1.0
source : 192.32.99.151
admin status : enabled
learning-mode : static
group : 234.10.10.0
source : 255.0.0.0
admin status : enabled
learning-mode : static
delete : N/A
disable : N/A
enable : N/A
```

Configuring multicast access control for an IGMP interface

Configure multicast access control for an IGMP interface or a VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Procedure steps

Configure multicast access control:

```
config ip igmp interface <ipaddr> access-control <name>
```

Variable definitions

Use the data in the following table to use the **`config ip igmp interface access-control`** and **`config ip vrf <vrfName> igmp interface access-control`** commands.

Variable	Value
create <*HostAddress*> <*HostMask*> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Creates an access control group entry for a specific IGMP interface. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host. • **`deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both`** indicates the action you want for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic.
delete <*HostAddress*> <*HostMask*>	Deletes the access control group entry for the specified IGMP interface.

Variable	Value
	• *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration.
info	Displays the settings for the access-control parameter.
ipaddr	Specifies the IP address of the selected interface.
mode <*HostAddress*> <*HostMask*> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Changes the access control group configuration. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration. • **`deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both`** indicates the action you want for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic.
name	Specifies the name of the access policy from 1–64 characters.

Configuring multicast access control for an IGMP Ethernet port

Configure multicast access control for an IGMP Ethernet port to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure steps

Configure multicast access control:

```
config ethernet <ports> ip igmp access-control <name>
```

Variable definitions

Use the data in the following table to use the **`config ethernet ip igmp access-control`** command.

Variable	Value
create <*HostAddress*> <*HostMask*> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Creates an access control group entry for a specific IGMP Ethernet port. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the host network. • **`deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both`** indicates the action you want for the specified IGMP Ethernet port. For example, if you specify deny-both, the interface denies both transmitted and received traffic.
delete <*HostAddress*> <*HostMask*>	Deletes the access control group entry for the specified IGMP interface. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration.
info	Displays the settings for the access-control parameter.
mode <*HostAddress*> <HostMask> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Changes the access control group configuration. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration. • **`deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both`** indicates the action you want for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic.

Variable	Value
name	Configures the name of the access policy from 1–64 characters.
ports	Specifies the Ethernet port using the convention {slot/port[-slot/port][,...]}.

Configuring multicast access control for a VLAN

Configure multicast access control for a selected IGMP VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure steps

Configure multicast access control:

```
config vlan <vid> ip igmp access-control <name>
```

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp access-control`** command.

Variable	Value
create <*HostAddress*> <*HostMask*> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Creates an access control group entry for a specified VLAN. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the host network. • **`deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both`** indicates the action you want for the specified VLAN. For example, if you specify deny-both, the VLAN denies both transmitted and received traffic.
delete <*HostAddress*> <*HostMask*>	Deletes the access control group entry for the specified VLAN.

Variable	Value
	• *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration.
info	Displays the settings for the access-control parameter.
mode <*HostAddress*> <*HostMask*> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Changes the access control group configuration. • *HostAddress* is the IP address of the host. • *HostMask* is the subnet mask used to determine the host or hosts covered by this configuration. • **`deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both`** indicates the action you want for the specified VLAN. For example, if you specify deny-both, the VLAN denies both transmitted and received traffic.
name	Configures the name of the access policy from 1–64 characters.
vid	Specifies a VLAN ID from 1–4092.

Configuring fast leave mode

Configure fast leave mode to specify if a port receives a leave message from a member of a group. Normal IGMP behavior is skipped. Fast leave mode provides one command that controls all IGMP fast leave enabled interfaces. You can use this global parameter to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces.

When a single user connects to an interface, you do not need to track if other users exist on the interface to perform the fast leave. In cases like this, you must change the mode to one-user.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces. It does not apply to IGAP interfaces, which ignore this mode.

You configure IGMP on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Procedure steps

1. View the current fast leave mode:

   ```
   show ip igmp info [vrf <value>] [vrfids <value>]
   ```

2. Configure fast leave mode:

   ```
   config ip igmp fast-leave-mode <one-user|multiple-user>
   ```

Variable definitions

Use the data in the following table to use the **`config ip igmp fast-leave-mode`** and **`config ip vrf <vrfName> igmp fast-leave-mode`** commands.

Variable	Value
one-user|multiple-user	**`one-user`** removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process. **`multiple-user`** removes from the group only the IGMP member who sent the leave message. Traffic is not stopped if other receivers exist on the interface port. Multiple-user is the default.

Example of configuring fast leave mode

Procedure steps

Change the mode to one-user.

```
ERS-8606:5# config ip igmp fast-leave-mode one-user
```

Configuring IGMP fast leave members on a VLAN

Configure IGMP fast leave members on a VLAN to specify fast leave capable ports.

Procedure steps

Configure fast leave members on a VLAN:

```
config vlan <vid> ip igmp fast-leave-members enable <ports>
```

Variable definitions

Use the data in the following table to use the **`config vlan ip igmp fast-leave-members`** command.

Variable	Value
disable *<ports>*	Prevents a port from receiving a leave message from a member of a group. Normal IGMP behavior is skipped.
enable *<ports>*	Enables members to join a fast leave group on a port on the VLAN. • *ports* is the port or list of ports that you want to join the fast leave group.
info	Displays information about the fast leave members of the VLAN.
vid	Specifies a VLAN ID from 1–4092.

Enabling IGMP L2 Querier globally

Enable L2 Querier to configure IGMP L2 Querier on a VLAN interface.

Procedure steps

1. Enable L2 Querier on a VLAN interface by using the following command:

   ```
   config vlan <vid> ip igmp snoop-querier enable
   ```

2. Disable L2 Querier on a VLAN interface by using the following command:

   ```
   config vlan <vid> ip igmp snoop-querier disable
   ```

Variable definitions

The following table describes variables that you enter in the **config vlan <vid> ip igmp snoop-querier <enable|disable>** command.

Variable	Value
<enable|disable>	Enables or disables snoop querier.
<vid>	Specifies the VLAQN ID in the range of 1 to 4094.

Configuring L2 snoop Querier address

Enable snoop Querier to configure an IP address on an interface.

Procedure steps

Configure snoop Querier address on a VLAN interface by using the following command:

```
config vlan <vid> ip igmp snoop-querier-addr <ipaddr>
```

Variable definitions

The following table describes variables that you enter in the **config vlan <vid> ip igmp snoop-querier-addr <ipaddr>** command.

Variable	Value
<ipaddr>	Specifies the IP address in the {a.b.c.d} format.
<vid>	Specifies the VLAQN ID in the range of 1 to 4094.

Configuring L2 Querier on an IP interface

Configure L2 Querier to enable L2 Querier on an IP interface.

Procedure steps

Configuring L2 Querier on an IP interface by using the following command:

```
config ip [vrf <name>] igmp interface <ipaddr> snoop-querier
<enable|disable>
```

Variable definitions

The following table describes variables that you enter in the **`config ip [vrf <name>] igmp interface <ipaddr> snoop-querier <enable|disable>`** command.

Variable	Value
<ipaddr>	Specifies the IP address in the {a.b.c.d} format.
<vid>	Specifies the VLAQN ID in the range of 1 to 4094.
vrf <name>	Specifies the VRF name.

Configuring L2 querier message source on a interface

Configure L2 querier message to set querier message source IP address.

Procedure steps

Configure querier message source on a interface by using the following command:

```
config ip [vrf <name>] igmp interface <ipaddr> snoop-querier-
addr <ipaddr>
```

Variable definitions

The following table describes variables that you enter in the **config ip [vrf <name>] igmp interface <ipaddr> snoop-querier-addr <ipaddr>** command.

Variable	Value
<ipaddr>	Specifies the IP address in the {a.b.c.d} format.
<vid>	Specifies the VLAQN ID in the range of 1 to 4094.
vrf <name>	Specifies the VRF name.

Viewing IP IGMP interface

Use the following procedure to view IP IGMP interface parameter details.

Procedure steps

View IP IGMP interface parameters by using the following command:

```
config ip [vrf <name>] igmp interface <ipaddr> info
```

Variable definitions

The following table describes variables that you enter in the **config ip [vrf <name>] igmp interface <ipaddr> info** command.

Variable	Value
<ipaddr>	Specifies the IP address in the {a.b.c.d} format.
vrf <name>	Specifies the VRF name.

Viewing IGMP snoop configuration result

Use the following procedure to view IGMP snoop configuration results.

Procedure steps

View IP IGMP interface parameters by using the following command:

```
show ip igmp snoop [vrf <WORD 0-64>] [vrfids <0-255>
```

Variable definitions

The following table describes variables that you enter in the **`show ip igmp snoop [vrf <value>] [vrfids <value>`** command.

Variable	Value
vrf <WORD 0–64>	Specifies the VRF name. The string length ranges from 0 to 64.
vrfids <0–255>	Specifies the VRF ID range.

Viewing IGMP configuration result

Use the following procedure to view IGMP configuration results.

Procedure steps

View IP IGMP interface parameters by using the following command:

```
show ip info igmp [<vid>][vrf <WORD 0-64>] [vrfids <0-255>
```

Variable definitions

The following table describes variables that you enter in the **`show ip info igmp [<vid>] [vrf <WORD 0-64>][vrfids <0-255>`** command.

Variable	Value
vrf <WORD 0–64>	Specifies the VRF name. The string length ranges from 0 to 64.
vrfids <0–255>	Specifies the VRF ID range.

Chapter 16: IGMP configuration using the ACLI

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.

For more information about IGMP configuration examples, see *PP8600 Technical Configuration Guide for IGMP Access Control*. You can find this document at http://www.avaya.com/support.

Prerequisites to IGMP configuration

- Complete one of the following tasks:
  - Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
  - Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Distance Vector Multicast Routing Protocol (DVMRP), Protocol Independent Multicast-Sparse Mode (PIM-SM), or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

Important:

To drop IGMPv2 control packets that do not use the router alert option, use the **`config ip igmp interface`** command and enable the router-alert parameter.

Navigation

- Job aid on page 356
- Configuring multicast stream limitation on an Ethernet port on page 358
- Configuring multicast stream limitation on a VLAN on page 360
- Configuring VLAN multicast stream limitation members on page 361
- Configuring multicast router discovery options on page 362
- Configuring IGMP static members on page 363
- Configuring SSM dynamic learning and range group on page 364

- Changing the SSM range group on page 366
- Configuring the SSM channel table on page 367
- Configuring multicast access control for an IGMP Ethernet port on page 369
- Configuring multicast access control for a VLAN on page 370
- Configuring fast leave mode on page 371
- Enabling fast leave mode on a port on page 373
- Configuring IGMP fast leave members on a VLAN on page 373
- Enabling IGMP L2 Querier globally on page 374
- Configuring L2 snoop Querier address on page 375
- Resetting L2 Querier on page 375
- Resetting snoop querier address on page 376
- Viewing IGMP snoop configuration result on page 376

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Important:

You must globally enable DVMRP or PIM multicasting on the switch for these commands to take effect.

Table 6: Job aid: Roadmap of IGMP ACLI commands

Command	Parameter
Privileged EXEC mode	
show ip igmp stream-limit interface	**<fastethernet|gigabitethernet>**
show ip igmp sys [vrf <WORD <0-64>] [vrfids <0-255>]	
show ip igmp snooping [vrf <WORD <0-64>] [vrfids <0-255>]	
Global Configuration mode	
ip igmp immediate-leave-mode	**<multiple-user|one-user>**

Command	Parameter
ip igmp ssm	dynamic-learning
	group-range <A.B.C.D/X>
ip igmp ssm-map	all
	<ip address> enable
	<IP address> <source IP address>
Interface Configuration mode	
ip igmp access-list <word> <host IP address>/<host mask address>	mode
	<deny-tx|deny-rx|deny-both| allow-only-tx|allow-only-rx| allow-only-both>
ip igmp immediate-leave	
ip igmp stream-limit	
ip igmp stream-limit-max-streams <0-65535>	
VLAN Interface Configuration mode	
ip igmp immediate-leave-members <ports>	
ip igmp mrdisc	maxadvertinterval <seconds>
	maxinitadvertinterval <seconds>
	maxinitadvertisements <integer>
	minadvertinterval <seconds>
	neighdeadinterval <seconds>
ip igmp static-group <group address> <to group address>	portList <static|blocked>
ip igmp stream-limit-group <ports>	enable
	max-streams <value>
ip igmp access-list <word> <host IP address>/<host mask address>	mode
	<deny-tx|deny-rx|deny-both| allow-only-tx|allow-only-rx| allow-only-both> ip igmp snoop-querier

Command	Parameter
ip igmp snoop-querier-addr <ipaddr>	
VRF Router Configuration mode	
ip igmp immediate-leave-mode	**<multiple-user|one-user>**
ip igmp ssm	**dynamic-learning**
	group-range <A.B.C.D/X>
ip igmp ssm-map	**all**
	<ip address> enable
	<IP address> <source IP address>

Configuring multicast stream limitation on an Ethernet port

Configure multicast stream limitation on an Ethernet port to limit the number of concurrent multicast streams on the port. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the port drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Prerequisites

- You must log on to the FastEthernet or GigabitEthernet Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

1. Enable multicast stream limitation:

   ```
   ip igmp stream-limit
   ```

2. Configure the maximum number of allowed streams:

```
ip igmp stream-limit-max-streams <0-65535>
```

3. Display multicast stream limitation information for the ports on a specific interface:

```
show ip igmp stream-limit interface [vrf Word<0-32>] [vrfids
Word<0-255>]
```

Variable definitions

Use the data in the following table to use the **ip igmp stream-limit-max-streams** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this port. The range is from 0–65535 and the default is 4. To use the default configuration, use the default option in the command: **default ip igmp stream-limit-max-streams**

Use the data in the following table to use the **show ip igmp stream-limit interface [vrf Word<0-32>] [vrfids Word<0-255>]** command.

Variable	Value
[vrf Word<0–32>]	Specifies the VRF name in the range of 0 to 32.
[vrfids Word<0–255>]	Specifies an VRF ID in the range of 0 to 255.

Example of configuring multicast stream limitation on an Ethernet port

Procedure steps

1. Enable multicast stream limitation on the Ethernet port.

```
ERS-8606:5(config-if)# ip igmp stream-limit
```

2. Set the maximum number of allowed streams to 8.

```
ERS-8606:5(config-if)# ip igmp stream-limit-max-streams 8
```

Job aid

The following table shows the field descriptions for this command.

Field	Description
INTERFACE	Indicates the interface IP address.
PORT	Indicates the port for the VLAN.
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

Configuring multicast stream limitation on a VLAN

Configure multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

You can configure the maximum number of streams independently. After a stream limit is met, the VLAN drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Enable multicast stream limitation:

   ```
   ip igmp stream-limit
   ```

2. Configure the maximum number of allowed streams:

   ```
   ip igmp stream-limit-max-streams <0-65535>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp`** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this port. The range is from 0–65535 and the default is 4. To use the default configuration, use the default option in the command: **default ip igmp stream-limit-max-streams**

Configuring VLAN multicast stream limitation members

Configure multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

Configure multicast stream limitation members on a VLAN:

```
ip igmp stream-limit-group <ports> enable max-streams <value>
```

Variable definitions

Use the data in the following table to use the **ip igmp stream-limit-group** command.

Variable	Value
max-streams <value>	Configures the maximum number of allowed streams for the specified ports on this VLAN. The range is from 0–65535 and the default is 4. To use the default configuration, use the default option in the command: **default ip igmp stream-limit-group <ports>**
<ports>	Specifies the port or port list in the format {slot/port[-slot/port][,...]}.

Example of configuring VLAN multicast stream limit members

Procedure steps

Enable multicast stream limitation on ports 1/3 to 1/8 and set the maximum allowed number of streams to 6 for this interface.

```
ERS-8606:5(config-if)# ip igmp stream-limit-group 1/3-1/8 max-streams 6
```

Configuring multicast router discovery options

Configure the multicast route discovery options to enable the automatic discovery of multicast-capable routers.

Important:

The Multicast Router Discovery (MRDISC) protocol is not supported on brouter ports.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Enable multicast router discovery:

   ```
   ip igmp mrdisc
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`ip igmp mrdisc`** command.

Variable	Value
maxadvertinterval <*seconds*>	Configures the maximum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration and reset the switch.

Variable	Value
	To set this option to the default value, use the **`default`** operator with the command. The default is 20.
maxinitadvertinterval <seconds>	Configures the maximum number (in seconds) between successive initial advertisements. For this change to take effect, you must save the configuration and reset the switch. To set this option to the default value, use the **`default`** operator with the command. The default is 2.
maxinitadvertisements <integer>	Configures the maximum number of initial multicast advertisements after initialization. For this change to take effect, you must save the configuration and reset the switch. To set this option to the default value, use the **`default`** operator with the command. The default is 3.
minadvertinterval <*seconds*>	Configures the minimum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration and reset the switch. To set this option to the default value, use the **`default`** operator with the command. The default is 15.
neighdeadinterval <*seconds*>	Configures the multicast router discovery dead interval—the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down. • *seconds* is a value from 1–59. The default is 60. To set this option to the default value, use the **`default`** operator with the command.

Configuring IGMP static members

Configure IGMP static members to add members to a snoop group. Some sources do not join a multicast group before transmitting a multicast stream. When this situation occurs and if no other group members joined in the VLAN, the data is flooded to all port members of the VLAN. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams are always forwarded to the multicast router within the VLAN, in addition to the ports configured for this static entry.

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

Configure interface static members:

```
ip igmp static-group <group address> <to group address>
[<portList>] <static|blocked>
```

Variable definitions

Use the data in the following table to use the **`ip igmp static-group`** command.

Variable	Value
<static|blocked>	Adds a static-member entry to the IGMP interface. • *value* is the port or list of ports to which you want to redirect the multicast stream for this multicast group. • **`static|blocked`** configures the route to static or blocked.
portList	Creates static members on the interface. Specifies the port or list of ports to which you want to redirect the multicast stream for this multicast group. Use the format {slot/port[-slot/port][,...]}. Use the **`no`** operator to later remove this configuration.
<group address> <to group address>	Indicates the IP address range {a.b.c.d[-w.x.y.z]} of the selected multicast group.

Configuring SSM dynamic learning and range group

Configure SSM dynamic learning and a range group to enable the IGMPv3 dynamic learning feature and to extend the default SSM range of 232/8 to include an IP multicast address. As new SSM channels are learned, they appear in the SSM channel table.

You configure IGMP on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- To define the range group, you must first disable PIM.
- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable SSM dynamic learning:

   ```
   ip igmp ssm dynamic-learning
   ```

2. Configure the range group:

   ```
   ip igmp ssm group-range <A.B.C.D/X>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp ssm`** command.

Variable	Value
A.B.C.D/X	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.1.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Example of configuring SSM dynamic learning and range group

Procedure steps

1. Define the SSM range group address (234.0.0.0) and mask (255.0.0.0).

```
ERS-8606:5(config)#ip igmp ssm group-range 234.0.0.0/255.0.0.0
```

2. Enable dynamic learning from IGMPv3 reports.

```
ERS-8606:5(config)# ip igmp ssm dynamic-learning
```

Changing the SSM range group

Change the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address.

Important:

This procedure reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), it also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap (BSR) is local.

You configure IGMP on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Disable PIM:

```
no ip pim enable
```

If you forget to disable PIM, the following error message appears:

```
Error: PIM is enabled in SSM mode, disable PIM
```

2. Delete each entry in the SSM channel table:

```
no ip igmp ssm-map [all] [IP address]
```

If you forget to delete the SSM channels, the following error message appears:

```
Error: SSM source group table not empty
```

3. Configure the new IP multicast group address:

   ```
   ip igmp ssm group-range <A.B.C.D/X>
   ```

4. Enable PIM:

   ```
   ip pim enable
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp ssm`** command.

Variable	Value
A.B.C.D/X	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.1.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Configuring the SSM channel table

Configure the SSM channel table to map groups to their sending source. SSM channels cannot conflict with static source groups and vice versa. After you configure an SSM channel or a static source group, the switch performs a consistency check to make sure no conflicts exist. You cannot map one group (G) to different sources for both a static source group and an SSM channel.

The consistency check applies to all SSM channel entries, even if they are disabled. If you disable an entry, it becomes inactive; you do not delete the entry and you can later reenable it.

After you disable an SSM channel, the Avaya Ethernet Routing Switch 8800/8600 stops multicast traffic from the specified source to the specified group. You can use this static setting as a security feature to block traffic from a certain source to a specific group.

You configure IGMP on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable the SSM channel table for all static entries:

   ```
   ip igmp ssm-map all
   ```

2. Enable a specific static entry:

   ```
   ip igmp ssm-map <IP address> enable
   ```

3. Create a static entry for a specific group:

   ```
   ip igmp ssm-map <IP address> <source IP address>
   ```

Variable definitions

Use the data in the following table to use the ip igmp ssm-map command.

Variable	Value
<IP address> enable	Enables the administrative state for a specific entry (group). This setting does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.
IP address <source IP address>	Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that is sends traffic to the group.

Example of configuring the SSM channel table

Procedure steps

1. Create an SSM channel table entry for the multicast group 234.0.1.0 and the source at 192.32.99.151.

   ```
   ERS-8606:5(config)# ip igmp ssm-map 234.0.1.0 192.32.99.151
   ```

2. Set the administrative state to enable all the static SSM channel table entries.

   ```
   ERS-8606:5(config)# ip igmp ssm-map all
   ```

Configuring multicast access control for an IGMP Ethernet port

Configure multicast access control for an IGMP Ethernet port to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI by selecting a specific port.

Procedure steps

1. Configure multicast access control:

   ```
   ip igmp access-list <word> <host IP address>/<host mask
   address> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-
   only-rx|allow-only-both>
   ```

2. Change an existing access list:

   ```
   ip igmp access-list <word> <host IP address>/<host mask
   address> mode <deny-tx|deny-rx|deny-both|allow-only-tx|
   allow-only-rx|allow-only-both>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp access-list`** command

Variable	Value
<host IP address> <host mask address>	Creates an access control group entry for a specific IGMP interface. • *host IP address* is the IP address of the host. • *host mask address* is the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic.
mode	Changes the access control group configuration.
word	Specifies the name of the access list from 1–64 characters.

Configuring multicast access control for a VLAN

Configure multicast access control for an IGMP VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Prerequisites

- You must log on to the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Configure multicast access control:

```
ip igmp access-list <word> <host IP address>/<host mask
address> <deny-tx|deny-rx|deny-both|allow-only-tx|allow-
only-rx|allow-only-both>
```

2. Change an existing access list:

```
ip igmp access-list <word> <host IP address> <host mask
address> mode <deny-tx|deny-rx|deny-both|allow-only-tx|
allow-only-rx|allow-only-both>
```

Variable definitions

Use the data in the following table to use the **`ip igmp access-list`** command.

Variable	Value
<host IP address> <host mask address>	Creates an access control group entry for a specific IGMP interface. • *host IP address* is the IP address of the host. • *host mask address* is the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic.
mode	Changes the access control group configuration.
word	Specifies the name of the access list from 1–64 characters.

Configuring fast leave mode

Configure fast (immediate) leave mode to specify if a port receives a leave message from a member of a group. Normal IGMP behavior is skipped. Fast leave mode provides one command that controls all IGMP fast leave enabled interfaces. You can use this global parameter to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces.

When a single user connects to an interface, you do not need to track if other users exist on the interface to perform the fast leave. In cases like this, you must change the mode to one-user.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces. It does not apply to Internet Group membership Authentication Protocol (IGAP) interfaces, which ignore this mode.

You configure IGMP on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. View the current fast leave mode:

   ```
   show ip igmp sys [vrf Word<0-16>] [vrfids Word<0-255>]
   ```

2. Configure fast leave mode:

   ```
   ip igmp immediate-leave-mode <multiple-user|one-user>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp immediate-leave-mode`** command.

Variable	Value
`multiple-user|one-user`	**`multiple-user`** removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This setting is the default. **`one-user`** removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.

Example of configuring fast leave mode

Procedure steps

Change the mode to one-user.

```
ERS-8606:5# ip igmp immediate-leave-mode one-user
```

Enabling fast leave mode on a port

Enable fast (immediate) leave mode to specify if a port receives a leave message from a member of a group. If you enable fast leave mode on a port, it uses the global fast leave mode configuration.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI by specifying a port or port list.

Procedure steps

Enable fast leave:

```
ip igmp immediate-leave
```

Configuring IGMP fast leave members on a VLAN

Configure IGMP fast leave members on a VLAN to specify fast leave capable ports.

Prerequisites

- You must log on the VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Enable fast leave on the VLAN:

   ```
   ip igmp immediate-leave
   ```

2. Configure fast leave members on a VLAN:

   ```
   ip igmp immediate-leave-members <ports>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp immediate-leave-members`** command.

Variable	Value
<ports>	Specifies the port or list of ports that you want to join the fast leave group.

Enabling IGMP L2 Querier globally

Enable L2 Querier to configure IGMP L2 Querier on a VLAN interface.

Prerequisites

- You must log on to the VLAN Configuration mode in the ACLI.

Procedure steps

1. Enable L2 Querier on a VLAN interface by using the following command:

```
ip igmp snoop-querier
```

2. Disable L2 Querier on a VLAN interface by using the following command:

```
no ip igmp snoop-querier
```

Configuring L2 snoop Querier address

Enable snoop Querier to configure an IP address on an interface.

Prerequisites

- You must log on to the VLAN Configuration mode in the ACLI.

Procedure steps

Configure L2 snoop Querier address on a VLAN interface by using the following command:

```
ip igmp snoop-querier-addr <ipaddr>
```

Resetting L2 Querier

Reset L2 Querier to set the L2 Querier to default value, that is disable.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI.

Procedure steps

Configuring L2 Querier on an IP interface by using the following command:

```
default ip igmp snoop-querier
```

Resetting snoop querier address

Reset snoop querier to set the snoop querier address to default value, that is 0.0.0.0.

Prerequisites

- You must log on to the Interface Configuration mode in the ACLI.

Procedure steps

Configure querier message source on a interface by using the following command:

```
default ip igmp snoop-querier-addr
```

Viewing IGMP snoop configuration result

Use the following procedure to view IGMP snoop configuration results.

Prerequisites

- You must log on to the Priv EXEC mode in the ACLI.

Procedure steps

View IP IGMP interface parameters by using the following command:

```
show ip igmp snooping [vrf <WORD 0-64>] [vrfids <0-255>
```

Variable definitions

The following table describes variables that you enter in the **`show ip igmp snooping [vrf <WORD 0-64>] [vrfids <0-255>`** command.

Variable	Value
vrf <WORD 0–64>	Specifies the VRF name. The string length ranges from 0 to 64.
vrfids <0–255>	Specifies the VRF ID range.

Chapter 17: PGM configuration using Enterprise Device Manager

Pragmatic General Multicast (PGM) provides reliable, duplicate-free delivery of data packets while reducing network congestion. PGM guarantees that receivers either can receive all data packets from transmissions and retransmissions, or can detect unrecoverable data packet loss. PGM is particularly well suited to push applications with relatively small information transfers such as stock and news updates. The Distance Vector Multicast Routing Protocol (DVMRP) is used between routers to exchange multicast routing information. For more information about PGM concepts and terminology, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP multicast Routing Protocols,* (NN46205-501).

The Avaya Ethernet Routing Switch 8800/8600 implements the network element portion of PGM. It supports the following PGM options:

- negative acknowledgement (NAK) list
- forward error correction (FEC)

Important:

The Avaya Ethernet Routing Switch 8800/8600 cannot serve as a designated local repairer (DLR) because DLRs require a large amount of buffering. Therefore, Enterprise Device Manager does not support the null negative acknowledgement (NNAK) parameters.

Prerequisites to PGM configuration

- You must configure the switch with Internet Group Management Protocol (IGMP) snoop or an IP multicast protocol such as DVMRP or Protocol Independent Multicast-Sparse Mode (PIM-SM). If PGM is configured without IP multicast enabled on a switch, PGM cannot run.
- Configure and enable IP multicast on the switch, particularly on the interfaces where PGM is required.

Navigation

- Enabling PGM globally on page 380
- Configuring VLANs with PGM on page 381

- Enabling PGM on an interface on page 382
- Editing PGM interface parameters on page 383
- Viewing PGM session parameters on page 384

Enabling PGM globally

Enable PGM globally to provide reliable, duplicate-free delivery of data packets while reducing network congestion.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PGM**.
3. Select **enabled**.
4. Click **Apply**.

Variable definitions

Use the data in the following table to configure the PGM Globals tab.

Variable	Value
Enable	Enables or disables PGM globally. The default is disabled.
State	Displays the current state (up or down) of PGM.
SessionLifeTime	Specifies the length of idle time (in seconds) after a session times out. Idle time is when the switch does not receive source path messages (SPM) from upstream routers. The default is 300 seconds.
NnakGenerate	If enabled, the router sends NNAK after it receives the redirect NCF. The default is enabled.
MaxReXmitStates	Configures the maximum number of retransmit state entries that the switch can create. Each entry uses a unique NAK sequence number. The default is 200 entries.
TotalReXmitStates	Displays the total number of retransmit state entries in the retransmit table.

Variable	Value
MaxSessions	Configures the maximum number of source path state sessions allowed on the switch. The default is 100 sessions.
TotalSessions	Displays the total number of source path state sessions in the PGM session entries table.
TotalReXmitStatesTimedOut	Displays the total number of retransmit state entries that were removed because they timed out.
TotalUniqueNaks	Displays the total number of unique NAKs received.
TotalUniqueParityNaks	Displays the total number of unique parity NAKs received.
MaxNakRate	Configures the maximum number of NAK transmission packets allowed for each second. The default is 100.

Configuring VLANs with PGM

Enable PGM on a VLAN to customize your PGM configuration.

Prerequisites

- Enable PGM globally.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select the VLAN ID that you want to configure with PGM.
4. Click **IP** from the menu bar.
5. Click the **PGM** tab.
6. Select **enabled**.
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the IP, VLAN dialog box, PGM tab.

Variable	Value
Enable	Enables or disables PGM on this interface. The default is disabled.
State	Indicates the current state (up or down) of PGM.
NakReXmitInterval	Specifies how long to wait for a NAK confirmation (NCF), in milliseconds, before retransmitting the NAK. The default is 1000 milliseconds.
MaxNakReXmitRate	Configures the maximum number of NAK retransmission packets allowed for each second. The default is 2.
NakRdataInterval	Specifies how long to wait for retransmitted data (RDATA), in milliseconds, after receiving an NCF. The default is 10 000 milliseconds.
NakEliminateInterval	Specifies the length of time (in milliseconds) that a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than NakRdataInterval. The default is 5000 milliseconds.

Enabling PGM on an interface

Enable PGM on an interface to customize your PGM configuration.

Prerequisites

- Enable PGM globally before you enable it on an interface.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.

3. Click **IP**.
4. Click the **PGM** tab.
5. Select **enabled**.
6. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Port dialog box, PGM tab.

Variable	Value
Enable	Enables or disables PGM on this interface. The default is disabled.
State	Indicates the current state (up or down) of PGM.
NakReXmitInterval	Specifies how long to wait for an NCF (in milliseconds) before retransmitting the NAK. The default is 1 000 milliseconds.
MaxNakReXmitRate	Configures the maximum number of NAK retransmission packets allowed for each second. The default is 2.
NakRdataInterval	Specifies how long to wait for RDATA (in milliseconds) after receiving an NCF. The default is 10 000 milliseconds.
NakEliminateInterval	Specifies the length of time (in milliseconds) that a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than NakRdataInterval. The default is 5 000 milliseconds.

Editing PGM interface parameters

Edit PGM interface parameters to modify the current configuration.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PGM**.
3. Click the **Interfaces** tab.

4. Click a field to edit it.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the PGM Interfaces tab.

Variable	Value
Cct	Displays the circuit number of the selected interface.
State	Displays the current state (up or down) of PGM.
Enable	Enables or disables PGM. The default is disabled.
NakReXmitInterval	Specifies how long to wait for an NCF (in milliseconds) before retransmitting the NAK. The default is 1000 milliseconds.
MaxNakReXmitRate	Configures the maximum number of NAK retransmission packets allowed for each second. The default is 2.
NakRdataInterval	Specifies how long to wait for RDATA (in milliseconds) after receiving an NCF. The default is 10 000 milliseconds.
NakEliminateInterval	Specifies the length of time (in milliseconds) that a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than NakRdataInterval. The default is 5000 milliseconds.

Viewing PGM session parameters

View PGM session parameters to view details of the current session.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **PGM**.
3. Click the **Session** tab.

Variable definitions

Use the data in the following table to use the PGM Session tab.

Variable	Value
Source	Displays the source IP address for this session.
Group	Displays the destination group address for this session.
SourcePort	Displays the source port for this session.
GlobalId	Displays the global ID for this session.
UpstreamAddress	Displays the IP address of the upstream interface for this session.
UpstreamIfCct	Displays the circuit number of the upstream interface for this session.
TrailEdgeSeq	Displays the trailing edge sequence of the transfer window.
LeadEdgeSeq	Displays the leading edge sequence of the transfer window.

Chapter 18: PGM configuration using the CLI

Pragmatic General Multicast (PGM) provides reliable, duplicate-free delivery of data packets while reducing network congestion. PGM guarantees that receivers either can either receive all data packets from transmissions and retransmissions, or detect unrecoverable data packet loss. PGM is particularly well suited to push applications with relatively small information transfers such as stock and news updates.

The Avaya Ethernet Routing Switch 8800/8600 implements the network element portion of PGM and supports the following PGM options:

- negative acknowledgement (NAK) list
- forward error correction (FEC)

Important:

The Avaya Ethernet Routing Switch 8800/8600 cannot serve as a designated local repairer (DLR) because DLRs require a large amount of buffering. Therefore, the CLI does not support the null negative acknowledgement (NNAK) parameters.

Prerequisites to PGM configuration

- You must configure the switch with IGMP snoop or an IP multicast protocol such as Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast-Spare Mode (PIM-SM). If PGM is configured without IP multicast enabled on a switch, PGM cannot operate.
- Configure and enable IP multicast on the switch, particularly on the interfaces where PGM is required.

PGM configuration navigation

- Job aid on page 388
- Configuring PGM globally on page 389
- Configuring PGM on an interface on page 390

• Configuring PGM on Ethernet ports on page 392
• Configuring PGM on a VLAN on page 393

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ethernet <ports> ip pgm`	`info`
	`max-nak-rdata-int <integer>`
	`max-nak-re-xmit-cnt <integer>`
	`nak-eliminate-int <integer>`
	`nak-re-xmit-int <integer>`
	`state <enable|disable>`
`config ip pgm`	`info`
	`max-rexmit-state <integer>`
	`max-sessions <integer>`
	`nnak-generate <enable| disable>`
	`state <enable|disable>`
	`session-life-time <integer>`
`config ip pgm interface <ipaddr>`	`info`
	`max-nak-rdata-int <integer>`
	`max-nak-re-xmit-cnt <integer>`
	`nak-eliminate-int <integer>`
	`nak-re-xmit-int <integer>`
	`state <enable|disable>`
`config vlan <vid> ip pgm`	`info`
	`max-nak-rdata-int <integer>`
	`max-nak-re-xmit-cnt <integer>`
	`nak-eliminate-int <integer>`

Command	Parameter
	nak-re-xmit-int <integer>
	state <enable|disable>
show ip pgm interface config	
show ip pgm interface error general	
show ip pgm interface error nak	
show ip pgm interface stat general	
show ip pgm interface stat nak	
show ip pgm interface stat parity	
show ip pgm retransmit	
show ip pgm session	
show ip pgm show-all [file <value>]	

Configuring PGM globally

Configure PGM globally on the switch to provide reliable, duplicate-free delivery of data packets while reducing network congestion.

Procedure steps

1. Enable PGM:

   ```
   config ip pgm state enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **config ip pgm** command.

Table 7: Variable definitons

Variable	Value
info	Displays current PGM settings on the switch.
max-rexmit-state *<integer>*	Configures the maximum number of retransmit state entries that the switch can create. Each entry uses a unique NAK sequence number. The range is 0–2147483647, and the default value is 200 entries.
max-sessions *<integer>*	Configures the maximum number of source path state sessions allowed on the switch. The range is 0–2147483647, and the default value is 100 sessions.
nnak-generate <enable|disable>	When enabled, the DLR that receives redirected NAKs, where it uses the retransmitted data (RDATA), sends an NNAK to the original source. The default is enable.
state <enable|disable>	Enables or disables PGM globally on the switch. The default is disable.
session-life-time *<integer>*	Specifies the length of idle time (in seconds) before a session times out. Idle time is when the switch does not receive source path messages (SPM) from the upstream routers. The range is 0–2147483647, and the default value is 300 seconds.

Configuring PGM on an interface

Configure PGM on an interface to customize your PGM configuration.

Prerequisites

- Enable PGM globally before you configure it on an interface.

Procedure steps

1. Enable PGM:

   ```
   config ip pgm interface <ipaddr> state enable
   ```

2. Configure the remaining parameters as required.
3. Verify the configuration on the interface:

```
show ip pgm interface config
```

Variable definitions

Use the data in the following table to use the **`config ip pgm interface`** command.

Variable	Value
info	Displays current PGM settings on the selected interface.
ipaddr	Indicates the IP address of the selected interface.
max-nak-rdata-int <*integer*>	Specifies how long to wait for RDATA, in milliseconds (ms), after receiving an NCF. The range is 1–2147483647, and the default value is 10 000 ms.
max-nak-re-xmit-cnt <*integer*>	Configures the maximum number of NAK retransmission packets allowed for each second. The range is 1–2147483647, and the default value is 2.
nak-eliminate-int <*integer*>	Specifies the length of time (in milliseconds) during which a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than **`max-nak-rdata-int`**. The range is 0–2147483647, and the default value is 5000 ms.
nak-re-xmit-int <*integer*>	Specifies how long to wait for a NAK confirmation (NCF), in milliseconds, before retransmitting the NAK. The range is 100–2147483647, and the default value is 1000 ms.
state <enable|disable>	Modifies the current state (enable or disable) of PGM on the selected interface. The default is disable.

Job aid

The following table describes the fields for this command.

Field	Description
CCT	Displays the circuit number of the selected interface.
ENABLE	Displays whether PGM is enabled or disabled on this interface.
STATE	Indicates the current state (up or down) of PGM.
NAK_RE_XMIT INTERVAL	Specifies how long to wait for an NCF (in milliseconds) before retransmitting the NAK. The default is 1000 milliseconds.

Field	Description
MAX_NAK_RE XMIT_COUNT	Displays the maximum number of NAK retransmission packets allowed for each second.
NAK_RDATA INTERVAL	Displays how long to wait for RDATA (in milliseconds) after receiving an NCF.
NAK_ELIMINATE INTERVAL	Displays the length of time (in milliseconds) during which a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than the NAK_RDATA INTERVAL.

Configuring PGM on Ethernet ports

Configure PGM at the port level to customize your PGM configuration.

Prerequisites

- Enable PGM globally before you configure it on Ethernet ports.

Procedure steps

1. Enable PGM on Ethernet ports:

   ```
   config ethernet <ports> ip pgm state enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config ethernet ip pgm`** command.

Variable	Value
info	Displays current PGM settings on the selected port.
max-nak-rdata-int *<integer>*	Specifies how long to wait for RDATA, in milliseconds (ms), after receiving an NCF. The range is 1–2147483647, and the default value is 10 000 ms.

Variable	Value
max-nak-re-xmit-cnt <*integer*>	Configures the maximum number of NAK retransmission packets allowed for each second. The range is 1–2147483647, and the default value is 2.
nak-eliminate-int <*integer*>	Specifies the length of time (in milliseconds) during which a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than **`max-nak-rdata-int`**. The range is 0–2147483647, and the default value is 5000 ms.
nak-re-xmit-int <*integer*>	Specifies how long to wait for a NAK confirmation (NCF), in milliseconds, before retransmitting the NAK. The range is 100–2147483647, and the default value is 1000 ms.
ports	Specifies the port using the convention {slot/port[-slot/port][,...]}.
state <enable|disable>	Modifies the current state (enable or disable) of PGM on the selected port. The default is disable.

Configuring PGM on a VLAN

Configure PGM on a VLAN to customize your PGM configuration.

Prerequisites

- Enable PGM globally before you configure it on a VLAN.

Procedure steps

1. Enable PGM on a VLAN:

   ```
   config vlan <vid> ip pgm state enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config vlan ip pgm`** command.

Variable	Value
info	Displays current PGM settings on the selected VLAN.
max-nak-rdata-int <*integer*>	Specifies how long to wait for RDATA, in milliseconds (ms), after receiving an NCF. The range is 1–2147483647, and the default value is 10 000 ms.
max-nak-re-xmit-cnt <*integer*>	Configures the maximum number of NAK retransmission packets allowed for each second. The range is 1–2147483647, and the default value is 2.
nak-eliminate-int <*integer*>	Specifies the length of time (in milliseconds) during which a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than **`max-nak-rdata-int`**. The range is 0–2147483647, and the default value is 5000 ms.
nak-re-xmit-int <*integer*>	Specifies how long to wait for a NAK confirmation (NCF), in milliseconds, before retransmitting the NAK. The range is 100–2147483647, and the default value is 1000 ms.
state <enable|disable>	Modifies the current state (enable or disable) of PGM on the selected VLAN. The default is disable.
vid	Specifies a VLAN ID from 1–4092.

Chapter 19: PGM configuration using the ACLI

Pragmatic General Multicast (PGM) provides reliable, duplicate-free delivery of data packets while reducing network congestion. PGM guarantees that receivers can either receive all data packets from transmissions and retransmissions, or detect unrecoverable data packet loss. PGM is particularly well suited to push applications with relatively small information transfers such as stock and news updates.

The Avaya Ethernet Routing Switch 8800/8600 implements the network element portion of PGM and supports the following PGM options:

- negative acknowledgement (NAK) list
- forward error correction (FEC)

Important:

The Avaya Ethernet Routing Switch 8800/8600 cannot serve as a designated local repairer (DLR) because DLRs require a large amount of buffering. Therefore, the ACLI does not support the null negative acknowledgement (NNAK) parameters.

Prerequisites to PGM configuration

- You must configure the switch with IGMP snoop or an IP multicast protocol such as Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast-Sparse Mode (PIM-SM). If PGM is configured without IP multicast enabled on a switch, PGM cannot run.
- Configure and enable IP multicast on the switch, particularly on the interfaces where PGM is required.

Navigation

- Job aid on page 396
- Configuring PGM globally on page 396
- Configuring PGM on a port or VLAN on page 398

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	
show ip pgm interface	**fastEthernet <slot/port>**
	gigabitEthernet <slot/port>
	vlan <id>
Global Configuration mode	
ip pgm	**enable**
	max-rexmit-state <integer>
	max-sessions <integer>
	nnak-generate <enable|disable>
	session-life-time <integer>
Interface Configuration mode	
ip pgm	**enable**
	max-nak-rdata-int <integer>
	max-nak-re-xmit-cnt <integer>
	nak-eliminate-int <integer>
	nak-re-xmit-int <integer>

Configuring PGM globally

Configure PGM globally on the switch to provide reliable, duplicate-free delivery of data packets while reducing network congestion.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable PGM:

   ```
   ip pgm enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`ip pgm`** command.

Variable	Value
max-rexmit-state <*integer*>	Configures the maximum number of retransmit state entries that the switch can create. Each entry uses a unique NAK sequence number. The default value is 200 entries. To set this option to the default value, use the **`default`** operator with the command.
max-sessions <*integer*>	Configures the maximum number of source path state sessions allowed on the switch. The default value is 100 sessions. To set this option to the default value, use the **`default`** operator with the command.
nnak-generate <enable| disable>	When enabled, the DLR that receives redirected NAKs, where it uses the retransmitted data (RDATA), sends an NNAK to the original source. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command. The default is enabled.
session-life-time <*integer*>	Specifies the length of idle time (in seconds) before a session times out. Idle time is when the switch does not receive SPMs from the upstream router. The default value is 300 seconds. To set this option to the default value, use the **`default`** operator with the command.

Configuring PGM on a port or VLAN

Configure PGM on an interface to customize your PGM configuration.

Prerequisites

- Enable PGM globally before you configure it on an interface.
- You must log on to the FastEthernet, GigabitEthernet, or VLAN Interface Configuration mode in the ACLI.

Procedure steps

1. Enable PGM:

   ```
   ip pgm enable
   ```

2. Configure the remaining parameters as required.
3. Verify the configuration on the interface:

   ```
   show ip pgm interface fastEthernet
   ```

   OR

   ```
   show ip pgm interface gigabitEthernet
   ```

   OR

   ```
   show ip pgm interface vlan
   ```

Variable definitions

Use the data in the following table to use the **`ip pgm`** command.

Variable	Value
max-nak-rdata-int *<integer>*	Specifies how long to wait for RDATA (in milliseconds) after receiving a NAK confirmation (NCF). The default value is 10000 milliseconds. To set this option to the default value, use the **`default`** operator with the command.

Variable	Value
max-nak-re-xmit-cnt <*integer*>	Configures the maximum number of NAK retransmission packets allowed for each second. The default value is 2 pps. To set this option to the default value, use the **`default`** operator with the command.
nak-eliminate-int <*integer*>	Specifies the length of time (in milliseconds) during which a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than **`max-nak-rdata-int`**. The default value is 5000 milliseconds. To set this option to the default value, use the **`default`** operator with the command.
nak-re-xmit-int <*integer*>	Specifies how long to wait for an NCF (in milliseconds), before retransmitting the NAK. The default value is 1000 milliseconds. To set this option to the default value, use the **`default`** operator with the command.

Job aid

The following table describes the fields for this command.

Field	Description
CCT	Displays the circuit number of the selected interface.
ENABLE	Displays whether PGM is enabled or disabled on this interface.
STATE	Indicates the current state (up or down) of PGM.
NAK_RE_XMIT INTERVAL	Specifies how long to wait for an NCF (in milliseconds) before retransmitting the NAK. The default is 1000 milliseconds.
MAX_NAK_RE XMIT_COUNT	Displays the maximum number of NAK retransmission packets allowed for each second.
NAK_RDATA INTERVAL	Displays how long to wait for RDATA (in milliseconds) after receiving a NCF.
NAK_ELIMINATE INTERVAL	Displays the length of time (in milliseconds) during which an NE eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than the NAK_RDATA INTERVAL.

Chapter 20: Route management using Enterprise Device Manager

Use the multicast tabs to view or edit interface configuration information for Layer 3 IP multicast protocols on the switch.

Navigation

- Viewing multicast route information on page 401
- Viewing multicast next-hop information on page 403
- Editing multicast interface information on page 404
- Editing static source groups on page 405
- Adding a new static source group on page 406
- Configuring a static IP multicast route on page 407
- Configuring IP multicast software forwarding on page 408
- Configuring the mroute stream limit on page 409
- Configuring the resource usage counter for multicast streams on page 410

Viewing multicast route information

View multicast route information for troubleshooting purposes.

Prerequisites

- To view route information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.

Variable definitions

Use the data in the following table to configure the Routes tab.

Variable	Value
Group	Displays the IP multicast group address for this entry that contains multicast routing information.
Source	Displays the network address that, when combined with the corresponding route SourceMask value, identifies the source containing multicast routing information.
SourceMask	Displays the network mask that, when combined with the corresponding route Source value, identifies the multicast source.
UpstreamNeighbor	Shows the address of the upstream neighbor from which the IP datagrams from these sources are received. The address is 0.0.0.0 if the network is local.
Interface	Displays the DVMRP interface, slot or port number, or VLAN ID where IP datagrams sent by these multicast sources to this multicast address are received.
ExpiryTime	Displays the amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the routing protocol through which the switch learned this route. Currently, the switch supports only the Distance Vector Multicast Routing Protocol (DVMRP).

Viewing multicast next-hop information

View all multicast next-hop information.

Prerequisites

- To view multicast next-hop information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Next Hops** tab.

Variable definitions

Use the data in the following table to configure the Next Hops tab.

Variable	Value
Group	Displays the IP multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
OutInterface	Displays the DVMRP interface slot or port number or VLAN ID for the outgoing interface for this next hop.
Address	Displays the address of the next hop specific to this entry. For most interfaces, it is identical to the next hop group. Non Broadcast Multiple Access (NBMA) interfaces, however, can use multiple next hop addresses out of a single outgoing interface.

Variable	Value
State	Displays whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. A value of forwarding indicates the information is currently used; pruned indicates it is not used.
ExpiryTime	Displays the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	Displays the minimum number of hops between this router and members of the IP multicast group reached through the next hop on this outgoing interface. IP multicast datagrams for the group that use a TTL less than this number of hops are not forwarded to the next hop.
Protocol	Displays the routing protocol through which the switch learned this next hop. Currently, the switch supports only the Distance Vector Multicast Routing Protocol (DVMRP).

Editing multicast interface information

Edit multicast interface information to customize your multicast configuration.

Prerequisites

- To configure multicast route information for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Interfaces** tab.
4. Click the field to edit it.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Multicast dialog box, Interfaces tab.

Variable	Value
Interface	Displays the slot or port number or VLAN ID for this entry.
Ttl	Configures the datagram time-to-live (TTL) threshold for the interface. The interface does not forward IP multicast datagrams with a TTL less than this threshold. The default value of 1 means that the interface forwards all multicast packets.
Protocol	Displays the routing protocol running on this interface. Currently only DVMRP is supported.

Editing static source groups

Configure static source-group entries in the DVMRP or Protocol Independent Multicast (PIM) multicast routing table. Neither DVMRP nor PIM can prune these entries from the distribution tree. In other words, even if no receivers exist in the group, the multicast stream for a static source-group entry stays active. The maximum number of static source groups must not exceed 1024.

Prerequisites

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - DVMRP
  - PIM-Sparse Mode (SM)
  - PIM-Source Specific Multicast (SSM)
- To configure static source groups on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.

3. Click the **Static Source Group** tab.
4. Edit the required information.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Static Source Group tab.

Variable	Value
GroupAddress	Configures the multicast group IP address for this static source-group entry.
SourceSubnet	Configures the multicast source address for this static source-group entry. How you configure the source address depends on the protocol and mode you use.
SrcSubnetMask	Configures the subnet mask of the source for this static source-group entry.

Adding a new static source group

Add a new static source group to create an entry that cannot be pruned from the distribution tree. An attempt to add a duplicate of an existing source-group entry results in an error message.

Prerequisites

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - DVMRP
  - PIM-SM
  - PIM-SSM
- To configure static source groups on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Static Source Group** tab.
4. Click **Insert**.
5. Complete the information in the dialog box.
6. Click **Insert**.

Variable definitions

Use the data in the following table to configure the Insert Static Source Group dialog box.

Variable	Value
GroupAddress	Configures the multicast group IP address for this static source-group entry.
SourceSubnet	Configures the multicast source address for this static source-group entry. How you configure the source address depends on the protocol and mode you use.
SrcSubnetMask	Configures the subnet mask of the source for this static source-group entry.

Configuring a static IP multicast route

Configure a static IP multicast route (mroute) to separate unicast and multicast traffic streams for Reverse Path Forwarding (RPF) calculation.

To configure an mroute on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IP**.

3. Click the **Static MRoutes** tab.
4. Click **Insert**.
5. Type the appropriate information for the static route.
6. Click **Insert**.

Variable definitions

Use the data in the following table to use the Static MRoutes tab.

Variable	Value
IpAddressType	Specifies the type of IP address for the static multicast route.
IpAddress	Specifies the IP address of the destination network for the mroute.
Mask	Specifies the network mask for the mroute.
RpfAddressType	Specifies the type of RPF IP address for the static multicast route.
RpfAddress	Specifies the IP address of the RPF neighbor. The RPF address is the same as the next hop in the unicast static IP route table.
Preference	Specifies a preference value from the range 1–255. The default is 1. Use the preference to assign an administrative distance of a route in the route table.
Enable	Specifies the status of the static multicast route. The default is enabled (selected).

Configuring IP multicast software forwarding

Configure IP multicast software forwarding to enable the Switch Fabric/Central Processor Unit (SF/CPU) to initially forward IP multicast data until a hardware record is created. The SF/CPU forwards the initial packets of a stream it receives and creates a corresponding hardware record for subsequent packets. The advantage of this feature is that it avoids initial data loss experienced by multicast applications and is most suited for low bandwidth.

The IP multicast software forwarding is a global system configuration feature that applies to all IP multicast-enabled interfaces and protocols. After you enable IP multicast software forwarding, the hardware still forwards IP multicast traffic; only initial data traffic is forwarded by the software. By default, the feature is disabled.

Important:

To avoid overloading the SF/CPU, Avaya recommends that you do not use the IP multicast software forwarding feature for video multicast applications.

Prerequisites

- To configure multicast software forwarding on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Mcast SW Forwarding** tab.
4. Select the **SWForwardingEnable** check box.
5. Click **Apply** to enable this feature.

Configuring the mroute stream limit

Limit the number of multicast streams to protect a SF/CPU from multicast data packet bursts generated by malicious applications, such as viruses that cause the SF/CPU to reach 100 percent utilization or that prevent the SF/CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a SF/CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

Important:

Avaya recommends that you configure more than 4000 multicast streams only with a SuperMezz module installed.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.

4. Click the **Mroute Stream Limit** tab.
5. Select the **StreamLimitEnable** checkbox.
6. Edit other fields as required.
7. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Mroute Stream Limit tab.

Variable	Value
StreamLimitEnable	Enables or disables mroute stream limit on the port. The default is disabled.
StreamLimit	Specifies the maximum number of multicast streams allowed to ingress to the SF/CPU through this port. The range is 1–32768, and the default is 1984.
StreamTimerCheck	Specifies the sampling period, in seconds, to check the number of multicast streams that ingressed the SF/CPU through this port. The range is 1–3600, and the default is 10.

Important:

Each user interface has unique terminology and naming conventions for parameters and values. For example, a parameter in EDM can appear in CLI with different spelling or syntax.

The following interface comparisons show examples of differences in terminology and syntax between identical parameters and values that you see when you configure and verify the mroute stream limit:

- EDM: displays the StreamTimerCheck parameter
- CLI: displays the MROUTE STR LIMIT TMR parameter

  and
- EDM: displays the StreamLimitEnable parameter with a value of enable or disable
- CLI: displays the ENABLE parameter with a value of true or false

Configuring the resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing your switch. After you set the counter thresholds for ingress and egress

records, if the record usage goes beyond the threshold, you receive a trap on the console, a logged message, or both.

Important:

If you do not set the thresholds, Enterprise Device Manager displays only the ingress and egress records currently in use.

Prerequisites

- To configure resource usage counters for a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Resource Usage** tab.
4. Type the required information.
5. Click **Apply**.

Variable definitions

Use the data in the following table to configure the Resource Usage tab.

Variable	Value
Egress Records In-Use	Displays the number of egress records (peps) traversing the switch.
Ingress Records In-Use	Displays the number of ingress records (source or group) traversing the switch.
Egress Threshold	Configures the egress threshold level (0–32767). A notification message is sent if this value is exceeded. The default value is 0.
Ingress Threshold	Configures the ingress threshold level (0–32767). A notification message is sent if this value is exceeded. The default value is 0.
Send Trap Only	Sends only trap notification messages after the threshold level is exceeded. Select disable if you select a different notification type. You can set only one notification type.

Variable	Value
SendTrapAndLog	Sends both trap and log notification messages after the threshold level is exceeded. Select disable if you select a different notification type.
LogMsgOnly	Sends only log notification messages after the threshold level is exceeded. Select disable if you select a different notification type.

Chapter 21: Route management using the CLI

With multicast route commands (mroute commands), you can configure and view IP multicast routing parameters on the Avaya Ethernet Routing Switch 8800/8600.

Route management navigation

- Job aid on page 413
- Displaying multicast routes on page 416
- Configuring a multicast route on an interface on page 416
- Configuring multicast stream limits on page 417
- Configuring multicast static source groups on page 419
- Configuring a static IP multicast route on page 421
- Configuring IP multicast software forwarding on page 422
- Configuring the resource usage counter for multicast streams on page 423

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ethernet <port> mroute-limit`	`info`
	`enable <true|false>`
	`max-allowed-streams <integer>`
	`max-allowed-streams-timer-check <integer>`
`config ip mroute`	`info`

Command	Parameter
`config ip mroute interface <ipaddr>`	`info`
	`ttl <ttl>`
`config ip mroute resource-usage`	`info`
	`egress-Threshold <integer>`
	`ingress-Threshold <integer>`
	`send-Trap-And-Log <enable|disable>`
	`trap-Msg-Only <enable|disable>`
	`log-Msg-Only <enable|disable>`
`config ip mroute static-source-group <GroupAddress>`	`info`
	`create <SourceSubnet> <SrcSubnetMask>`
	`delete <SourceSubnet> <SrcSubnetMask>`
`config ip static-mroute`	`create <ipaddr/mask> rpf <value> [preference <value>]`
	`delete <ipaddr/mask> rpf <value>`
	`disable <ipaddr/mask> rpf <value>`
	`enable <ipaddr/mask> rpf <value>`
	`info`
	`preference <value> <ipaddr/mask> rpf <value>`
`config ip vrf <WORD 0-64> mroute`	`info`
`config ip vrf <WORD 0-64> mroute interface <ipaddr>`	`info`
	`ttl <ttl>`
`config ip vrf <WORD 0-64> mroute resource-usage`	`info`
	`egress-Threshold <integer>`
	`ingress-Threshold <integer>`

Command	Parameter
	send-Trap-And-Log <enable| disable>
	trap-Msg-Only <enable| disable>
	log-Msg-Only <enable|disable>
config ip vrf <WORD 0-64> mroute static-source-group <GroupAddress>	info
	create <SourceSubnet> <SrcSubnetMask>
	delete <SourceSubnet> <SrcSubnetMask>
config ip vrf <WORD 0-64> static-mroute	create <ipaddr/mask> rpf <value> [preference <value>]
	delete <ipaddr/mask> rpf <value>
	disable <ipaddr/mask> rpf <value>
	enable <ipaddr/mask> rpf <value>
	info
	preference <value> <ipaddr/mask> rpf <value>
config sys mcast-software-forwarding	info
	disable
	enable
config sys vrf <WORD 0-64> mcast-software-forwarding	info
	disable
	enable
show ip mroute-hw group-prune-state	
show ip mroute-hw group-trace	
show ip mroute-hw resource-usage [vrf <WORD 0-64>] [vrfids <0-255>]	

Command	Parameter
show ip mroute rpf <ipaddr> [vrf <WORD 0-64>] [vrfids <0-255>]	
show ip mroute show-all [file <value>]	
show ports info mroute-limit	
show sys mcast-software-forwarding	

Displaying multicast routes

Display multicast routes to view the current configuration.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

Display routes:

```
config ip mroute info
```

Configuring a multicast route on an interface

Configure a multicast route on an interface to customize the time-to-live (TTL) value for the route.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

Configure a route:

```
config ip mroute interface <ipaddr> ttl <ttl>
```

Variable definitions

Use the data in the following table to use the **config ip mroute interface** and **config ip vrf <vrfName> mroute interface** commands.

Variable	Value
info	Displays information about the multicast route interface.
ipaddr	Indicates the IP address of the selected interface.
ttl *<ttl>*	Configures the datagram time-to-live (TTL) threshold for the interface. The interface does not forward IP multicast datagrams with a TTL less than this threshold. The range (in seconds) is 1–255. The default is 1.

Configuring multicast stream limits

Limit the number of multicast streams to protect a Switch Fabric/Central Processor Unit (SF/CPU) from multicast data packet bursts generated by malicious applications, such as viruses that cause the SF/CPU to reach 100 percent utilization or that prevent the SF/CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to an SF/CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

You can enable or disable the mroute stream limit for the entire device or for an individual port when the switch is operating. If you enable the mroute stream limit for the device and for an individual port, only the periodic check is performed for that port.

Important:

Avaya recommends that you configure more than 4000 multicast streams only if you use a SuperMezz module.

Procedure steps

1. Enable stream limits:

   ```
   config ethernet <port> mroute-limit enable true
   ```

2. Configure the remaining parameters as required.
3. Show the mroute stream limit configuration:

```
show ports info mroute-limit
```

Variable definitions

Use the data in the following table to use the **config ethernet mroute-limit** command.

Variable	Value
enable <true|false>	Enables or disables the feature on the specified port. By default it is disabled.
info	Displays the current configuration.
max-allowed-streams *<integer>*	Configures the maximum number of streams for the specified port. The port is shut down if the number of streams exceeds this limit. The value is a number between 1–32768. The default value is 1984 streams.
max-allowed-streams-timer-check *<integer>*	Configures the sampling interval, which is used to check if the number of ingress multicast streams to the SF/CPU is under a configured limit or if the port needs to shut down. The range is between 1–3600. The default value is 10 seconds.
port	Specifies the port or range of ports in slot or port notation.

Job aid

The following message appears while shutting down the port due to excessive multicast streams:

```
Shutdown port <port> due to excessive multicast streams <# of streams
ingressed>; Configured limit max streams <configured limit> in
<configured sampling interval> sec. Please disable and re-enable the
port.
```

The following table shows the field descriptions for this command.

Field	Description
PORT	Indicates the port number.
MROUTE STR LIMIT	Indicates the maximum number of multicast streams that can ingress to the SF/CPU through this port.

Field	Description
MROUTE STR LIMIT TIMER	Indicates the sampling period (in seconds) to check the number of multicast streams that ingressed the SF/CPU through this port.
ENABLE	Indicates the status of the mroute stream limit on the port.

Important:

Each user interface has unique terminology and naming conventions for parameters and values. For example, a parameter in EDM can appear in CLI with different spelling or syntax.

The following interface comparisons show examples of differences in terminology and syntax between identical parameters and values that you see when you configure and verify the mroute stream limit:

- EDM: displays the StreamTimerCheck parameter
- CLI: displays the MROUTE STR LIMIT TMR parameter

  and
- EDM: displays the StreamLimitEnable parameter with a value of enable or disable
- CLI: displays the ENABLE parameter with a value of true or false

Configuring multicast static source groups

Configure static source-group entries in the Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) multicast routing table. Neither DVMRP nor PIM can prune these entries from the distribution tree. In other words, even if no receivers exist in the group, the multicast stream for a static source-group entry stays active. The maximum number of static source groups must not exceed 1024.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **`config ip`** with **`config ip vrf <vrfName>`** in the following procedure.

Prerequisites

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - DVMRP
  - PIM-Sparse Mode (SM)

- PIM-Source Specific Multicast (SSM)

Procedure steps

Configure a static source-group entry:

```
config ip mroute static-source-group <GroupAddress> create
<SourceSubnet> <SrcSubnetMask>
```

Variable definitions

Use the data in the following table to use the **`config ip mroute static-source-group`** and **`config ip vrf <vrfName> mroute static-source-group`** commands.

Variable	Value
create <*SourceSubnet*> <*SrcSubnetMask*>	Creates a new static multicast source-group entry. You cannot create duplicate groups. • *SourceSubnet* is the multicast source address for this static source-group entry. How you configure the source address depends on the protocol and mode you use. • *SrcSubnetMask* is the subnet mask of the source for this static source-group entry.
delete <*SourceSubnet*> <*SrcSubnetMask*>	Deletes the source-group entry from the static source-group table. • *SourceSubnet* is the multicast source address for this static source-group entry. • *SrcSubnetMask* is the subnet mask of the source for this static source-group entry.
info	Displays information about the source-group entry.
GroupAddress	Specifies the IP address of the multicast group.

Example of configuring multicast static source groups

Procedure steps

Create a static source group for two multicast groups: 224.32.2.1 and 226.50.2.2.

Tip:

The static mroute for group 224.32.2.1 is for a source subnet 10.10.10.0/24.

The static mroute for group 226.50.2.2 is for the host 20.20.20.100/32.

```
ERS-8606:5#config ip mroute static-source-group 224.32.2.1
ERS-8606:5/config/ip/mroute/static-source-group/224.32.2.1#create
10.10.10.0 255.255.255.0
ERS-8606:5/config/ip/mroute/static-source-group/224.32.2.1#config ip mroute
static-source-group 226.50.2.2
ERS-8606:5/config/ip/mroute/static-source-group/226.50.2.2# create
20.20.20.100 255.255.255.255
```

Configuring a static IP multicast route

Configure a static IP multicast route (mroute) to separate unicast and multicast traffic streams for Reverse Path Forwarding (RPF) calculation.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Create an mroute:

   ```
   config ip static-mroute create <ipaddr/mask> rpf <value>
   [preference <value>]
   ```

2. Enable an mroute:

   ```
   config ip static-mroute enable <ipaddr/mask> rpf <value>
   ```

3. Disable an mroute:

   ```
   config ip static-mroute disable <ipaddr/mask> rpf <value>
   ```

4. Modify the administrative distance of an mroute:

   ```
   config ip static-mroute preference <value> <ipaddr/mask> rpf
   <value>
   ```

5. Delete an mroute:

   ```
   config ip static-mroute delete <ipaddr/mask> rpf <value>
   ```

6. View the mroute table configuration:

   ```
   config ip static-mroute info
   ```

7. View the mroute to reach a specific RPF address:

   ```
   show ip mroute rpf <ipaddr> [vrf <value>] [vrfids <value>]
   ```

Important:

Because the switch uses two types of static IP routes, one for unicast routes and one for multicast routes, the command **`show ip route info`** does not provide the correct RPF information to a source or RP. Avaya recommends that you use the **`show ip mroute rpf`** command.

Variable definitions

Use the data in the following table to use the **`config ip static-mroute`** and **`config ip vrf <vrfName> static-mroute`** commands.

Variable	Value
ipaddr/mask	Specifies the IP address and the network mask for the mroute. Use the format a.b.c.d/x or a.b.c.d/x.x.x.x
preference <value>	Specifies a preference value from the range 1–255. The default is 1.
rpf <value>	Specifies the IP address of the RPF neighbor.
[vrf <value>] [vrfids <value>]	Specifies a VRF name or a range of VRF IDs to include in the show command results.

Configuring IP multicast software forwarding

Configure the IP multicast software forwarding feature so the SF/CPU initially forwards IP multicast data until a hardware record is created. The SF/CPU forwards the initial packets of a stream it receives and creates a corresponding hardware record for subsequent packets. The advantage of this feature is that it avoids initial data loss experienced by multicast applications and is most suited for low bandwidth.

The IP multicast software forwarding is a global system configuration feature that applies to all IP multicast-enabled interfaces and protocols. After you enable IP multicast software forwarding, the hardware still forwards IP multicast traffic; the software only forwards initial data traffic. By default, the feature is disabled.

Important:

To avoid overloading the SF/CPU, Avaya recommends that you do not use the IP multicast software forwarding feature for video multicast applications.

You configure multicast software forwarding on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Enable software forwarding:

   ```
   config sys mcast-software-forwarding enable
   ```

2. Show the software forwarding configuration:

   ```
   show sys mcast-software-forwarding [vrf <value>] [vrfids
   <value>]
   ```

Variable definitions

Use the data in the following table to use the **config sys mcast-software-forwarding** and **config sys vrf <vrfName> mcast-software-forwarding** commands.

Variable	Value
disable	Disables IP multicast software forwarding. This setting is the default.
enable	Enables IP multicast software forwarding. The default is disable.
info	Displays the current configuration information.

Configuring the resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing your switch. After you set the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive a trap on the console, a logged message, or both.

If you do not set the thresholds, the CLI displays only the ingress and egress records currently in use.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Configure the resource usage counter options:

   ```
   config ip mroute resource-usage
   ```

2. Enable traps and log messages:

   ```
   config ip mroute resource-usage send-Trap-And-Log enable
   ```

Variable definitions

Use the data in the following table to use the **`config ip mroute resource-usage`** and **`config ip vrf <vrfName> mroute resource-usage`** commands.

Variable	Value
egress-Threshold <*integer*>	Configures the egress record threshold (S,G). A notification message is sent if this value is exceeded. • *integer* is a value between 0–32767. The default it 0.
ingress-Threshold <*integer*>	Configures the ingress record threshold (peps). A notification message is sent if this value is exceeded. • *integer* is a value between 0–32767. The default is 0.
info	Displays the current configuration for the record usage counter.
send-Trap-And-Log <enable|disable>	Configures the notification method for sending both a trap message and a log message after the threshold level is exceeded. The default is disable. **Important:** You can only set one notification type.
trap-Msg-Only <enable|disable>	Configures the notification method for sending only a trap message after the threshold level is exceeded. The default is disable.
log-Msg-Only <enable|disable>	Configures the notification method for sending only a log message after the threshold level is exceeded. The default is disable.

Example of configuring the resource usage counter

Procedure steps

1. Set the egress threshold to 200.

```
ERS_8606:5# config ip mroute resource-usage egress-Threshold 200
```

2. Set the ingress threshold to 100.

```
ERS_8606:5/config/ip/mroute/resource-usage# ingress-Threshold 100
```

3. Enable the log message notification method.

```
ERS_8606:5/config/ip/mroute/resource-usage# log-Msg-Only enable
```

Chapter 22: Route management using the ACLI

With multicast route commands, you can configure and view IP multicast routing parameters on the Avaya Ethernet Routing Switch 8800/8600.

Route management navigation

- Job aid on page 427
- Configuring a multicast route on an interface on page 429
- Configuring multicast stream limits on page 430
- Configuring multicast static source groups on page 432
- Configuring a static IP multicast route on page 433
- Configuring IP multicast software forwarding on page 434
- Configuring the resource usage counter for multicast streams on page 435

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Privileged EXEC mode	
show ip mroute interface <fastethernet| gigabitethernet>	**<slot/port>**
show ip static-mroute	**ip <A.B.C.D>**
	rpf <A.B.C.D>
	vrf <Word/0-32>

Command	Parameter
	vrfids <0-255>
show multicast software-forwarding [vrf Word<0-16>] [vrfids Word<0-255>]	
Global Configuration mode	
ip mroute	stream-limit
ip mroute interface	<ipaddr>
	ttl <1-255>
ip mroute resource-usage	egress-threshold <0-32767>
	ingress-threshold <0-32767>
	log-msg
	trap-msg
ip static-mroute	<A.B.C.D/0-32> rpf <A.B.C.D> [preference <1-255>] [enable]
	<A.B.C.D/0-32> preference <1-255>
ip mroute static-source-group	<A.B.C.D>
	<A.B.C.D/X>
multicast software-forwarding	
Interface Configuration mode	
ip mroute	max-allowed-streams
	max-allowed-streams-timer-check
	port
	stream-limit
VRF Router Configuration mode	
ip mroute	stream-limit
ip mroute interface	<ipaddr>
	ttl <1-255>
ip mroute resource-usage	egress-threshold <0-32767>
	ingress-threshold <0-32767>
	log-msg

Command	Parameter
	trap-msg
ip mroute static-source-group	**<A.B.C.D>**
	<A.B.C.D/X>
ip static-mroute	**create <A.B.C.D/0-32> rpf <A.B.C.D> [preference <1-255>] [enable]**
	<A.B.C.D/0-32> preference <1-255>
multicast software-forwarding	

Configuring a multicast route on an interface

Configure a multicast route on an interface to customize the time-to-live (TTL) value for the route.

You configure a multicast route on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Configure a route:

```
ip mroute interface <ipaddr> ttl <1-255>
```

Variable definitions

Use the data in the following table to use the **ip mroute interface** command.

Variable	Value
ipaddr	Indicates the IP address of the selected interface.
ttl <1–255>	Configures the TTL threshold for the multicast route interface. The range, in seconds, is 1–255. To set this option to the default value, use the **default** operator with the command. The default is 1.

Configuring multicast stream limits

Limit the number of multicast streams to protect a Switch Fabric/Central Processor Unit (SF/CPU) from multicast data packet bursts generated by malicious applications, such as viruses that cause the SF/CPU to reach 100 percent utilization or that prevent the SF/CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to an SF/CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

You can enable or disable the mroute stream limit for the entire device or for an individual ports when the switch is operating. If you enable the mroute stream limit for the device and for an individual port, the switch performs only the periodic check for that port.

Important:

Avaya recommends that you configure more than 4000 multicast streams only if you use a SuperMezz module.

You configure stream limits on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Procedure steps

1. Log on to the Global Configuration mode.
2. Enable stream limitation globally:

   ```
   ip mroute stream-limit
   ```

3. Log on to the FastEthernet, GigabitEthernet, or VLAN Interface Configuration mode.
4. Enable stream limits:

   ```
   ip mroute stream-limit
   ```

5. For FastEthernet or Gigabit Ethernet interfaces, configure the remaining parameters as required.
6. Show the mroute stream limit configuration:

```
show ip mroute interface <fastethernet|gigabitethernet>
[<slot/port>]
```

Variable definitions

Use the data in the following table to use the **ip mroute** command.

Variable	Value
max-allowed-streams *<integer>*	Configures the maximum number of streams on the specified port. The port is shut down if the number of streams exceeds this limit. The value is a number between 1–32768. The default value is 1984 streams. To set this option to the default value, use the **default** operator with the command.
max-allowed-streams-timer-check *<integer>*	Configures the sampling interval, which is used to check if the number of ingress multicast streams to the SF/CPU is under a configured limit or if the port needs to shut down. The range is between 1–3600. The default value is 10 seconds. To set this option to the default value, use the **default** operator with the command.
port	Specifies the port or range of ports in slot or port notation. Use the **no** operator to later remove this configuration.

Job aid

The following message appears while shutting down the port due to excessive multicast streams:

```
Shutdown port <port> due to excessive multicast streams <# of streams
ingressed>; Configured limit max streams <configured limit> in
<configured sampling interval> sec. Please disable and re-enable the
port.
```

The following table shows the field descriptions for the **show ip mroute interface** command.

Field	Description
PORT	Indicates the port number.
MROUTE STR LIMIT	Indicates the maximum number of multicast streams that can ingress to the SF/CPU through this port.

Field	Description
MROUTE STR LIMIT TIMER	Indicates the sampling period (in seconds) to check number of multicast streams that use ingressed the SF/CPU through this port.
ENABLE	Indicates the status of the mroute stream limit on the port.

Configuring multicast static source groups

Configure static source-group entries in the Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) multicast routing table. Neither DVMRP nor PIM can prune these entries from the distribution tree. Even if no receivers exist in the group, the multicast stream for a static source-group entry remains active. The maximum number of static source groups must not exceed 1024.

You configure multicast static source groups on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - DVMRP
  - PIM-Sparse Mode (SM)
  - PIM-Source Specific Multicast (SSM)
- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Configure a static source-group entry:

```
ip mroute static-source-group <A.B.C.D> <A.B.C.D/X>
```

Variable definitions

Use the data in the following table to use the **`ip mroute static-source-group`** command.

Variable	Value
A.B.C.D	Specifies the IP address of the multicast group. Use the **no** operator to remove this configuration.
A.B.C.D/X	Specifies the multicast source IP address and subnet mask for the static source group entry. You cannot create duplicate groups. How you configure the source address depends on the protocol and mode you use. Use the **no** operator to remove this configuration.

Example of configuring multicast static source groups

Procedure steps

Create a static source group for two multicast groups: 224.32.2.1 and 226.50.2.2. The static mroute for group 224.32.2.1 is for a source subnet 10.10.10.0/24. The static mroute for group 226.50.2.2 is for the host 20.20.20.100/32.

```
ERS-8606:5(config)#ip mroute static-source-group 224.32.2.1 10.10.10.0/24
ERS-8606:5(config)# ip mroute static-source-group 226.50.2.2 20.20.20.100/32
```

Configuring a static IP multicast route

Configure a static IP multicast route (mroute) to separate unicast and multicast traffic streams for Reverse Path Forwarding (RPF) calculation.

You configure an mroute on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode.

Procedure steps

1. Create and enable an mroute:

   ```
   ip static-mroute <A.B.C.D/0-32> rpf <A.B.C.D> [preference <1-255>] [enable]
   ```

2. Disable an mroute:

```
no ip static-mroute <A.B.C.D/0-32> rpf <A.B.C.D> enable
```

3. Delete an mroute:

```
no ip static-mroute <A.B.C.D/0-32> rpf <A.B.C.D>
```

4. View the mroute to reach a specific RPF address:

```
show ip static-mroute [ip <A.B.C.D>] [rpf <A.B.C.D>] [vrf
<Word 0-64>] [vrfids <0-255>]
```

Variable definitions

Use the data in the following table to use the **`ip static-mroute`** command.

Variable	Value
1–255	Specifies a preference value from the range 1–255. The default value is 1. To set this option to the default value, use the **`default`** operator with the command.
A.B.C.D	Specifies the IP address of the RPF neighbor or the static route.
A.B.C.D/0-32	Specifies the IP address and the network mask for the mroute.
[vrf <WORD 0–64>][vrfids <0–255>]	Specifies a VRF name or a range of VRF IDs to include in the show command results.

Configuring IP multicast software forwarding

Configure the IP multicast software forwarding feature so the SF/CPU initially forwards IP multicast data until a hardware record is created. The SF/CPU forwards the initial packets of a stream it receives and creates a corresponding hardware record for subsequent packets. The advantage of this feature is that it avoids initial data loss experienced by multicast applications and is most suited for low bandwidth.

The IP multicast software forwarding is a global system configuration feature that applies to all IP multicast-enabled interfaces and protocols. After you enable IP multicast software forwarding, the hardware still forwards IP multicast traffic; the software only forwards initial data traffic. By default, the feature is disabled.

Important:

To avoid overloading the SF/CPU, Avaya recommends that you do not use the IP multicast software forwarding feature for video multicast applications.

You configure multicast software forwarding on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable software forwarding:

   ```
   multicast software-forwarding
   ```

2. Show the software forwarding configuration:

   ```
   show multicast software-forwarding [vrf Word<0-16>] [vrfids
   Word<0-255>]
   ```

Configuring the resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing your switch. After you configure the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive a trap on the console, a logged message, or both.

If you do not set the thresholds, the ACLI displays only the ingress and egress records currently in use.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Configure the resource usage counter options:

```
ip mroute resource-usage egress-threshold <0-32767> ingress-
threshold <0-32767>
```

2. Enable traps and log messages:

```
ip mroute resource-usage log-msg trap-msg
```

Variable definitions

Use the data in the following table to use the **`ip mroute resource-usage`** command.

Variable	Value
egress-threshold *<0–32767>*	Configures the egress record threshold (S,G). A notification message is sent if this value is exceeded. • *integer* is a value between 0–32767. To set this option to the default value, use the **`default`** operator with the command. The default is 0.
ingress-threshold *<0–32767>*	Configures the ingress record threshold (peps). A notification message is sent if this value is exceeded. • *integer* is a value between 0–32767. To set this option to the default value, use the **`default`** operator with the command. The default is 0.
log-msg	Configures the notification method for sending only a log message after the threshold level is exceeded. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.
trap-msg	Configures the notification method for sending only a trap message after the threshold level is exceeded. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command. The default is disabled.

Example of configuring the resource usage counter

Procedure steps

1. Set the egress threshold to 200.

```
ERS-8606:5(config)#ip mroute resource-usage egress-threshold 200
```

2. Set the ingress threshold to 100.

```
ERS-8606:5(config)# ip mroute resource-usage ingress-threshold 100
```

3. Enable the log message notification method.

```
ERS-8606:5(config)# ip mroute resource-usage log-msg
```

Chapter 23: Multicast flow distribution over MLT using Enterprise Device Manager

Multicast flow distribution over MultiLink Trunking (MLT) provides a mechanism for distributing multicast streams over a multilink trunk. With MLT, you can distribute the load on different ports of the multilink trunk and (whenever possible) achieve an even distribution of the streams.

To configure multicast flow distribution over MLT, you must enable it globally and for each multilink trunk.

Multicast flow distribution over MLT procedures

This task flow shows you the sequence of procedures you perform to configure multicast flow distribution over MLT by using the Device Manager Enterprise. To link to a procedure, select Multicast flow distribution over MLT navigation on page 440.

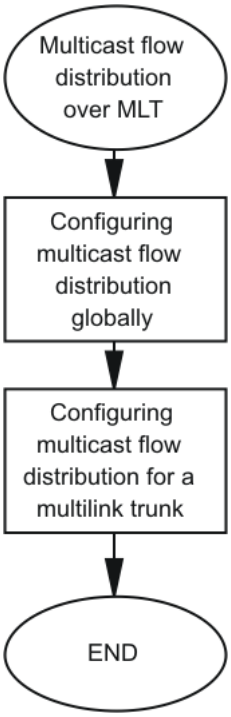

Figure 22: Multicast flow distribution over MLT procedures using the Device Manager Enterprise

Multicast flow distribution over MLT navigation

- Configuring multicast flow distribution globally on page 440
- Configuring multicast flow distribution for a multilink trunk on page 441

Configuring multicast flow distribution globally

Configure multicast flow distribution globally to distribute multicast streams over a multilink trunk.

All MLT distribution for unicast or multicast traffic occurs at the ingress port, and therefore the port style make-up of the MLT itself has no affect on this operation. If the ingress port is a R, RS, or 8800 module port and the egress MLT is in the same VLAN (Layer 2 flow) distribution occurs only if IGMP snoop is enabled.

For inter-VLAN or Layer 3 flows, distribution always occurs for ingress IPMC traffic. In this configuration, you must enable an IPMC Layer 3 routing protocol, or configure static IPMC routing. Distribution does not occur for a pure Layer 2 multicast traffic (multicast destination MAC only) that has no Layer 3 IPMC address. If IGMP snoop is enabled for a particular VLAN/Brouter port, the switch accepts Layer 3 multicast protocol configuration on that VLAN or brouter port.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **Edit**.
2. Click **Chassis**.
3. Click the **Mcast MLT Distribution** tab.
4. Select **Enable** in the check box.
5. Optionally, type a group IP address in the **GrpMask** box.
6. Optionally, type a source IP address in the **SrcMask** box.
7. Optionally, select the **RedistributionEnable** check box.
8. Click **Apply**.

   After you select enable or disable, or redistribution enable or disable, the following information message appears:

   ```
   Multicast distribution over MLT must be enabled/disabled on
   both sides of the MLT, otherwise loops or traffic
   ```

```
interruption or traffic interruption occur. Do you want to
continue?
```

9. Click **Yes** to continue or **No** to cancel the operation.

   If you click Yes, multicast flow distribution over MLT is globally configured.

Variable definitions

Use the data in the following table to configure the Mcast MLT Distribution tab.

Variable	Value
Enable	Globally adctivates multicast flow distribution. The default is disabled.
GrpMask	Configures a group mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255.
SrcMask	Configures a source mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255.
RedistributionEnable	Enables or disables the multicast MLT redistribution feature. The default is disabled.

Configuring multicast flow distribution for a multilink trunk

Enable multicast flow distribution for each multilink trunk to customize your configuration. Distribute the load on different ports of the multilink trunk.

Prerequisites

- Enable multicast flow distribution globally.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **MLT/LACP**.
3. Click the **MultiLink/LACP Trunks** tab.

4. Click **Insert**.
5. In the **Id** box, type the ID number for the multilink trunk.
6. In the **PortType** box, select the port type.
7. In the **Name** box, type a name for the multilink trunk port.
8. In the **PortMembers** box, click the ellipsis **[...]** button, select the ports, and then click **Ok**.
9. In the **VlanIds** box, click the **[...]** button, select a VLAN, and then click **Ok**.
10. In the **MltType** box, select the multilink trunk type.
11. If you choose **splitMLT**, in the SmltId box, type the SMLT ID (a value from 1 to 256).

    Important:

    You must pair the SMLT ID on each aggregation switch. The SMLT ID is the identification number that the internal spanning tree (IST) uses to determine the split multilink trunk to which to send information. This number is identified between the two aggregation switches.
12. In the **MulticastDistribution** box, select **enable** to activate multicast flow distribution.
13. Select the **NtStgEnable** check box to ensure that the spanning tree group is in Avaya mode.
14. In the **Aggregatable** box, select **disable**.
15. In the **AggMinLink** box, type the number of active AggMinLink.
16. Click **Insert**.

Variable definitions

Use the data in the following table to configure the MultiLink/LACP Trunks tab.

Variable	Value
Id	Configures a value that uniquely identifies the multilink trunk associated with this entry.
PortType	Configures the port type as either access or trunk port. The default is access.
Name	Specifies the name of the multilink trunk.
PortMembers	Specifies the ports assigned to the multilink trunk.
VlanIds	Specifies the VLANs to which the ports belong.

Variable	Value
MltType	Configures the type of multilink trunk. The options include: • normalMLT • istMLT • splitMLT The default is normalMLT.
SmltId	Specifies the split multilink trunk ID.
MulticastDistribution	Enables or disables multicast flow distribution. The default is disable.
NtStgEnable	Indicates whether this spanning tree group (STG) is operating in Avaya mode or in Cisco mode. True is Avaya mode, and false is Cisco mode. The default setting is true.
Aggregatable	Enables or disables the link aggregation on a multilink trunk. The default is disable.
AggMinLink	Specifies the number of active links that goes below the configured minimum-link number. The entire LAG is declared down.

Chapter 24: Multicast flow distribution over MLT using the CLI

Multicast flow distribution over Multilink Trunking (MLT) provides a mechanism for distributing multicast streams over an multilink trunk. You can distribute the load on different ports of the multilink trunk and (whenever possible) achieve an even distribution of the streams.

To configure multicast flow distribution over MLT, you must enable it globally and for each multilink trunk. For more information about MLT, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Link Aggregation, MLT and SMLT, NN46205-518*.

Multicast flow distribution over MLT procedures

This task flow shows you the sequence of procedures you perform to configure multicast flow distribution over MLT. To link to a procedure, go toMulticast flow distribution over MLT navigation on page 446.

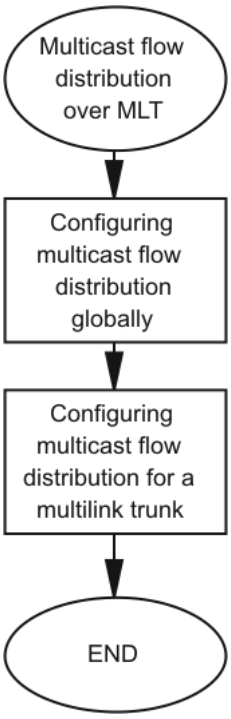

Figure 23: Multicast flow distribution over MLT procedures

Multicast flow distribution over MLT navigation

- Job aid on page 446
- Configuring multicast flow distribution globally on page 446
- Configuring multicast flow distribution for a multilink trunk on page 448

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
config mlt <mltid> mcast-distribution	**disable**
	enable
config sys mcast-mlt-distribution	**info**
	disable
	enable
	grp-mask <grp-mask>
	redistribution <enable| disable>
	src-mask <src-mask>

Configuring multicast flow distribution globally

Configure multicast flow distribution globally to distribute multicast streams over a multilink trunk.

All MLT distribution for unicast or multicast traffic occurs at the ingress port, and therefore the port style make-up of the MLT itself has no affect on this operation. If the ingress port is a R, RS, or 8800 module port and the egress MLT is in the same VLAN (Layer 2 flow) distribution occurs only if IGMP snoop is enabled.

For inter-VLAN or Layer 3 flows, distribution always occurs for ingress IPMC traffic. In this configuration, you must enable an IPMC Layer 3 routing protocol, or configure static IPMC routing. Distribution does not occur for a pure Layer 2 multicast traffic (multicast destination

MAC only) that has no Layer 3 IPMC address. If IGMP snoop is enabled for a particular VLAN/ Brouter port, the switch accepts Layer 3 multicast protocol configuration on that VLAN or brouter port.

Procedure steps

1. Enable multicast flow distribution:

   ```
   config sys mcast-mlt-distribution enable
   ```

2. Configure the remaining parameters as required.

Variable definitions

Use the data in the following table to use the **`config sys mcast-mlt-distribution`** command.

Variable	Value
disable	Disables multicast flow distribution over MLT. This setting is the default.
enable	Enables multicast flow distribution over MLT globally.
grp-mask <*grp-mask*>	Specifies a group mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255.
info	Displays the current configuration information.
redistribution <enable| disable>	Enables or disables the multicast MLT redistribution feature. The default is disabled.
src-mask <*src-mask*>	Specifies a source mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255. **Important:** Ensure that the mask values for grp-mask and src-mask are contiguous.

Example of configuring multicast flow distribution

Procedure steps

1. Enable multicast flow distribution over MLT globally.

```
ERS-8606:5# config sys mcast-mlt-distribution enable
```

2. Set the mask for the group so that it takes into account the last byte of the group address.

```
ERS-8606:5# config sys mcast-mlt-distribution grp-mask 0.0.0.255
```

3. Set the mask for the source so that it takes into account the last two bytes of the source IP address.

```
ERS-8606:5# config sys mcast-mlt-distribution src-mask 0.0.255.255
```

4. Enable redistribution to allow streams to redistribute if changes occur in the multilink trunk.

```
ERS-8606:5# config sys mcast-mlt-distribution redistribution enable
ERS-8606:5/config/sys/mcast-mlt-distribution# redistribution enable
```

```
IpmMltMcastDistributionGrpConsistencyCheck: Enable multicast
redistribution over MLT disrupt traffic for existing streams
on the MLT during redistribution
```

```
ERS-8606:5/config/sys/mcast-mlt-distribution# enable
```

Configuring multicast flow distribution for a multilink trunk

Enable multicast flow distribution for each multilink trunk to customize your configuration. Distribute the load on different ports of the multilink trunk.

Procedure steps

Enable multicast flow distribution:

```
config mlt <mltid> mcast-distribution enable
```

Variable definitions

Use the data in the following table to use the **`config mlt mcast-distribution`** command.

Variable	Value
enable	Enables multicast flow distribution on the specified multilink trunk.
disable	Disables multicast flow distribution on the specified multilink trunk. This setting is the default.
mltid	Specifies the MLT ID, in the range of 1–32.

Chapter 25: Multicast flow distribution over MLT using the ACLI

Multicast flow distribution over MultiLink Trunking (MLT) provides a mechanism for distributing multicast streams over an multilink trunk. You can distribute the load on different ports of the multilink trunk and (whenever possible) to achieve an even distribution of the streams.

To configure multicast flow distribution over MLT, you must enable it globally and for each multilink trunk. For more information about MLT, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Link Aggregation and Multilink Trunking,* (NN46205-506).

Multicast flow distribution over MLT procedures

This task flow shows you the sequence of procedures you perform to configure multicast flow distribution over MLT. To link to a procedure, go to Multicast flow distribution over MLT navigation on page 452.

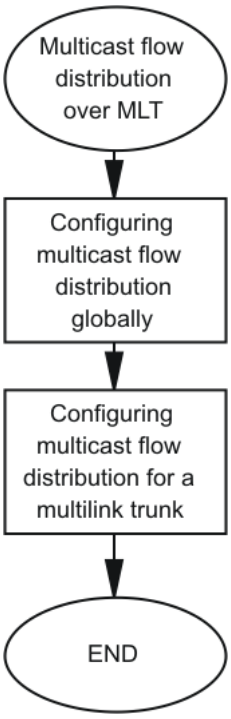

Figure 24: Multicast flow distribution over MLT procedures

Multicast flow distribution over MLT navigation

- Job aid on page 452
- Configuring multicast flow distribution globally on page 452
- Configuring multicast flow distribution for a multilink trunk on page 454

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Global Configuration and MLT Interface Configuration modes	
multicast mlt-distribution	**grp-mask <grp-mask>**
	redistribution
	src-mask <src-mask>

Configuring multicast flow distribution globally

Configure multicast flow distribution globally to distribute multicast streams over a multilink trunk.

All MLT distribution for unicast or multicast traffic occurs at the ingress port, and therefore the port style make-up of the MLT itself has no affect on this operation. If the ingress port is a R, RS, or 8800 module port and the egress MLT is in the same VLAN (Layer 2 flow) distribution occurs only if IGMP snoop is enabled.

For inter-VLAN or Layer 3 flows, distribution always occurs for ingress IPMC traffic. In this configuration, you must enable an IPMC Layer 3 routing protocol, or configure static IPMC routing. Distribution does not occur for a pure Layer 2 multicast traffic (multicast destination MAC only) that has no Layer 3 IPMC address. If IGMP snoop is enabled for a particular VLAN/ Brouter port, the switch accepts Layer 3 multicast protocol configuration on that VLAN or brouter port.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Configure the optional parameters as required:

   ```
   multicast mlt-distribution [grp-mask <grp-mask>]
   [redistribution] [src-mask <src-mask>]
   ```

2. Enable multicast flow distribution:

   ```
   multicast mlt-distribution
   ```

Variable definitions

Use the data in the following table to use the **`multicast mlt-distribution`** command.

Variable	Value
grp-mask *<grp-mask>*	Specifies a group mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255. To set this option to the default value, use the **`default`** operator with the command. You must disable MLT distribution before you change the group mask.
redistribution	Enables the multicast MLT redistribution feature. The default is disabled. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command.
src-mask *<src-mask>*	Specifies a source mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255. **Important:** Ensure that the mask values for grp-mask and src-mask are contiguous. To set this option to the default value, use the **`default`** operator with the command.

Example of configuring multicast flow distribution

Procedure steps

1. Set the mask for the group so that it takes into account the last byte of the group address.

```
ERS-8606:5(config)# multicast mlt-distribution grp-mask 0.0.0.255
```

2. Set the mask for the source so that it takes into account the last two bytes of the source IP address.

```
ERS-8606:5(config)# multicast mlt-distribution src-mask 0.0.255.255
```

3. Enable redistribution to allow streams to redistribute if changes occur in the multilink trunk.

```
ERS-8606:5(config)# multicast mlt-distribution redistribution
```

```
IpmMltMcastDistributionGrpConsistencyCheck: Enable multicast
redistribution over MLT disrupt traffic for existing streams
on the MLT during redistribution
```

4. Enable multicast flow distribution over MLT globally.

```
ERS-8606:5(config)#multicast mlt-distribution
```

Configuring multicast flow distribution for a multilink trunk

Enable multicast flow distribution for each multilink trunk to customize your configuration. Distribute the load on different ports of the multilink trunk.

Prerequisites

- You must log on to the MLT Interface Configuration mode in the ACLI.

Procedure steps

1. Configure the optional parameters as required:

```
multicast mlt-distribution [grp-mask <grp-mask>]
[redistribution] [src-mask <src-mask>]
```

2. Enable multicast flow distribution:

```
multicast mlt-distribution
```

Variable definitions

Use the data in the following table to use the **`multicast mlt-distribution`** command.

Variable	Value
grp-mask <*grp-mask*>	Specifies a group mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255. To set this option to the default value, use the **`default`** operator with the command. You must disable MLT distribution before you change the group mask.
redistribution	Enables the multicast MLT redistribution feature. The default is disabled. Use the **`no`** operator to later remove this configuration. To set this option to the default value, use the **`default`** operator with the command.
src-mask <*src-mask*>	Specifies a source mask to use when the switch distributes multicast traffic over a multilink trunk. The default is 255.255.255.255. **Important:** Ensure that the mask values for grp-mask and src-mask are contiguous. To set this option to the default value, use the **`default`** operator with the command.

Chapter 26: MVR configuration using Enterprise Device Manager

This chapter describes how to configure the Multicast VLAN Registration (MVR) Protocol using Enterprise Device Manager.

For conceptual information about MVR, see Multicast VLAN Registration Protocol on page 73.

Navigation

- Enabling MVR globally on page 457
- Enabling MVR or proxy on a VLAN on page 458
- Adding a receiver VLAN to the MVR VLAN on page 459
- Disabling MVR on page 460
- Viewing the MVR group on page 460
- Viewing the MVR VLAN information on page 461

Enabling MVR globally

Enable MVR to save bandwidth and maintain the performance of the multicast router when multicast groups are located in different VLANs.

Prerequisites

- Disable Protocol Independent Multicast (PIM).

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.

3. Click **IGMP**. The IGMP dialog box appears with the Global tab displayed.
4. Click **Global** tab.
5. In the **Mvr** box, select **enable**.
6. Click **Apply**.

Enabling MVR or proxy on a VLAN

Enable MVR or MVR Proxy on a VLAN to designate the MVR VLAN for the switch.

Prerequisites

- Disable Protocol Independent Multicast (PIM).
- Enable MVR globally.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.
3. Click **IGMP**.

   The IGMP dialog box appears with the MVR Vlans tab displayed.
4. Click the **MVR Vlans** tab.
5. In the **MvrVlanProxy** box of the desired VLAN, click and select **true**.

Variable definitions

Use the data in this table to help you configure the Mvr Vlans tab.

Variable	Value
MvrVlanId	Specifies the Mcast-VLAN ID.
MvrVlanProxy	Specifies if the proxy working mode is enabled or disabled for Mcast-VLAN.
SourcePort	Specifies the source port for the Mcast-VLAN.

Variable	Value
SourceAddress	Specifies the querier address for the source for Mcast-VLAN.

Adding a receiver VLAN to the MVR VLAN

Add a receiver VLAN to the MVR VLAN to bind the VLAN.

Prerequisites

- Enable MVR globally.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.
3. Click **IGMP**.

   The IGMP dialog box appears with the Global tab displayed.
4. Select the **MVR Receivers** tab.
5. Click **Insert**.

   The Insert MVR Receivers dialog box appears
6. In the **VlanId** box, type the ID number of the desired VLAN.
7. Click **Insert**.

Variable definitions

Use the data in the following table to help you configure the Insert VLAN dialog box fields.

Variable	Value
VlanId	Specifies the ID number of the specified VLAN.

Disabling MVR

Disable MVR so that you can enable PIM.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.
3. Click **IGMP**.

   The IGMP dialog box appears with the Global tab displayed.
4. Click **Global** tab.
5. In the **Mvr** box, select **disable**.
6. Click **Apply**.

Viewing the MVR group

View the MVR information to verify the configuration.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.
3. Click **IGMP**.
4. Click the **MVR Groups** tab.
5. Review the information.

Variable definitions

Use the data in the following table to help you configure the MVR Groups tab fields.

Variable	Value
GroupAddress	Specifies the IP address of each multicast group.
VlanId	Specifies the ID numbers of the receiver VLANs.
IGMPVersion	Specifies the version of IGMP.
LivingTime	Specifies how long the group exist.

Viewing the MVR VLAN information

View the MVR VLAN information to learn the MVR VLAN ID and to verify the configuration.

Procedure steps

1. In the Device Physical View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration** > **IP**.
3. Click **IGMP**.

   The IGMP dialog box appears with the Global tab displayed.
4. Click the **MVR Vlans** tab.
5. View the information.

Variable definitions

Use the data in the following table to help you configure MVR Vlans tab fields.

Variable	Value
MvrVlanId	Specifies the ID number of the MVR VLAN.
MvrVlanProxy	Specifies whether MVR VLAN Proxy is enabled.
SourcePort	The port that is connected to the multicast router.
SourceAddress	The IP address of the source address.

Chapter 27: MVR configuration using the CLI

This chapter provides the procedures to use to configure the Multicast VLAN Registration (MVR) Protocol by using the Command Line Interface (CLI).

For conceptual information about MVR, see Multicast VLAN Registration Protocol on page 73.

Navigation

- Enabling MVR globally on page 464
- Configuring MVR on a VRF on page 465
- Enabling MVR or MVR Proxy on a VLAN on page 466
- Adding a receiver VLAN to the MVR VLAN on page 467
- Viewing MVR group information on page 468
- Viewing multicast VLAN information on page 469

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Table 8: Job aid: Roadmap of MVR CLI commands

Command	Parameter
`config ip igmp multicast-vlan-registration`	`<enable|disable>`
`config ip vrf <vrfname> igmp multicast-vlan-registration`	`<create|delete>`
`config vlan <vid> ip igmp multicast-vlan-registration`	`<enable|disable>`

Command	Parameter
config vlan <vid> ip igmp multicast-vlan-registration-proxy	**<enable|disable>**
config ip igmp multicast-vlan-registration	**add receiver-vlan <vid>**
	info
	remove receiver-vlan <vid>
show ip igmp multicast-vlan-registration group [vlan <value> [address <value>] [vrf <value>] [vrfids <value>]	
show ip igmp multicast-vlan-registration vlan [vrf <value>] [vrfids <value]	
show ip igmp multicast-vlan-registration vlan [vrf <value>] [vrfids <value]	

Enabling MVR globally

Enable MVR to save bandwidth and maintain the performance of the multicast router when multicast groups are located in different VLANs.

Prerequisites

- Disable Protocol Independent Multicast (PIM).

Procedure steps

1. Enable MVR by using the following command:

   ```
   config ip igmp multicast-vlan-registration enable
   ```

2. Disable MVR by using the following command:

   ```
   config ip igmp multicast-vlan-registration disable
   ```

Variable definitions

The following table describes variables that you enter in the **`config ip igmp multicast-vlan-registration enable`** command.

Variable	Value
<enable|disable>	Enables or disables MVR globally.

The following table describes optional parameters the you enter after the **`config ip igmp multicast-vlan-registration`** command.

Variable	Value
add receiver-vlan <value>	Adds receiver VLAN. **`vid`** is the VLAN ID in the range of 1 to 4094.
info	Shows current level parameter settings and next level directories.
remove receiver-vlan <value>	Deletes receiver VLAN. `vid` is the VLAN ID in the range of 1 to 4094.

Configuring MVR on a VRF

Use this procedure to create or delete a MVR instance on a VRF.

Prerequisites

- Disable Protocol Independent Multicast (PIM).
- Enable MVR globally.

Procedure steps

1. Create a MVR instance on a VRF:

   ```
   config ip vrf <vrfname> igmp multicast-vlan-registration
   create
   ```

2. Remove a MVR instance on a VRF:

   ```
   config ip vrf <vrfname> igmp multicast-vlan-registration
   delete
   ```

Note:

MVR for Global Router is created by default. Default MVR is disabled when you create a new MVR instance.

Variable definitions

Variable	Value
add receiver-vlan <vid>	Adds receiver VLAN. vid is the VLAN ID, range of 1 to 4094.
create	Creates an MVR instance.
delete	Removes an MVR instance.
disable	Disables MVR.
enable	Enables MVR.
info	Shows the current level parameter settings and next level directories.
remove receiver-vlan <vid>	Deletes a receiver VLAN. vid is the VLAN ID, range of 1 to 4094.

Enabling MVR or MVR Proxy on a VLAN

Enable MVR or MVR Proxy on a VLAN to designate the MVR VLAN for the switch.

Prerequisites

- Disable Protocol Independent Multicast (PIM).
- Enable MVR globally.

Procedure steps

1. Enable a VLAN as the MVR VLAN for the switch by using the following command:

   ```
   config vlan <vid> ip igmp multicast-vlan-registration enable
   ```

2. Enable a VLAN as the MVR Proxy VLAN for the switch by using the following command:

```
config vlan <vid> ip igmp multicast-vlan-registration-proxy
enable
```

Variable definitions

The following table describes variables that you enter in the **config vlan <vid> ip igmp multicast-vlan-registration enable** and **config vlan <vid> ip igmp multicast-vlan-registration-proxy enable** commands.

Variable	Value
<enable|disable>	Enables or disables the MVR VLAN.
<vid>	Specifies the MVR VLAN ID in the range of 1 to 4094.

Adding a receiver VLAN to the MVR VLAN

Add a receiver VLAN to the MVR VLAN to bind the VLAN.

Prerequisites

- Disable Protocol Independent Multicast (PIM).
- Enable MVR globally, when adding the first receiver VLAN. You can subsequently add more receiver VLANs even when MVR is disabled.
- Enable IGMP snoop on a VLAN.

Procedure steps

1. Bind a receiver VLAN to the MVR VLAN by using the following command:

   ```
   config ip igmp multicast-vlan-registration add receiver-vlan
   <value>
   ```

2. Display the global MVR configuration by using the following command:

   ```
   config ip igmp multicast-vlan-registration info
   ```

Variable definitions

The following table describes variables that you enter in the **config ip igmp multicast-vlan-registration add receiver-vlan <value>** command.

Variable	Value
<value>	Specifies the VLAN ID in the range of 1 to 4094.

The following table describes optional parameters the you enter after the **config ip igmp multicast-vlan-registration** command.

Variable	Value
disable	Disables MVR.
enable	Enables MVR.
remove receiver-vlan <value>	Unbinds a VLAN from the MVR VLAN. *value* is the ID of the VLAN and is in the range of 1 to 4094.

Viewing MVR group information

View MVR group information to verify the configuration and troubleshoot.

Procedure steps

View the MVR group by using the following command:

```
show ip igmp multicast-vlan-registration group [vlan <value>]
[address <value] [vrf <value>] [vrfids <value>]
```

Variable definitions

The following table describes variables that you enter in the **show ip igmp multicast-vlan-registration group [vlan <value>] [address <value] [vrf <WORD 0-64>] [vrfids <0-255>]** command.

Variable	Value
[vlan <value>]	Specifies the VLAN to which the group belongs. *value* is the VLAN ID in the range of 1 to 4094. This is only applicable for a receiver VLAN. The VLAN ID must be a receiver VLAN.
[address <value>]	Specifies the IP address of a specific group. *value* is the IP address in the format of A.B.C.D.
vrf <WORD 0–64>	Specifies the VRF name. The string length ranges from 0 to 64.
vrfids <0–255>	Specifies the VRF ID range.

The following table describes optional parameters the you enter after the **`show ip igmp multicast-vlan-registration`** command.

Variable	Value
info [vrf <WORD 0–64>] [vrfids <0–255>]	Displays current level parameter settings and next level directories. • **`vrf <WORD 0–64>`** is the VRF name. The string length ranges from 0 to 64. • **`vrfids <0–255>`** is the VRF ID range.
group [vlan <value>] [address <value>] [vrf <WORD 0–64>] [vrfids <0–255>]	Displays the Multicast VLAN registration group information. • **`vrf <WORD 0–64>`** is the VRF name. The string length ranges from 0 to 64. • **`vrfids <0–255>`** is the VRF ID range.
vlan [vrf <WORD 0–64>] [vrfids <0–255>]	Displays the Multicast VLAN registration information. • **`vrf <WORD 0–64>`** is the VRF name. The string length ranges from 0 to 64. • **`vrfids <0–255>`** is the VRF ID range.

Viewing multicast VLAN information

View multicast VLAN information to see the VLANS that are bound to the MVR VLAN.

Procedure steps

Display the multicast VLAN information by using the following command:

```
show ip igmp multicast-vlan-registration vlan [vrf <WORD 0-64>]
[vrfids <0-255>]
```

Variable definitions

The following table describes variables that you enter in the **`show ip igmp multicast-vlan-registration vlan [vrf <value>] [vrfids <value>]`** commands.

Variable	Value
vrf <WORD 0–64>	Specifies the VRF name. The string length ranges from 0 to 64.
vrfids <0–255>	Specifies the VRF ID range.

The following table describes optional parameters that you enter after the **`show ip igmp multicast-vlan-registration`** command.

Variable	Value
group [vlan <value>] [address <value>] [vrf <WORD 0–64>] [vrfids <0–255>]	Displays the Multicast VLAN registration information. • **`vrf <WORD 0-64>`** is the VRF name. The string length ranges from 0 to 64. • **`vrfids <0-255>`** is the VRF ID range.
info [vrf <WORD 0–64>] [vrfids <0–255>]	Displays current level parameter settings and next level directories. • **`vrf <WORD 0-64>`** is the VRF name. The string length ranges from 0 to 64. • **`vrfids <0-255>`** is the VRF ID range.

Chapter 28: MVR configuration using the ACLI

This chapter provides the procedures to use to configure the Multicast VLAN Registration (MVR) Protocol by using the Avaya command line interface (ACLI).

For conceptual information about MVR, see Multicast VLAN Registration Protocol on page 73.

Prerequisites

- Disable Protocol Independent Multicast (PIM)

Navigation

- Job aid on page 471
- Enabling MVR globally on page 472
- Disabling MVR globally on page 473
- Configuring MVR on VRF using ACLI on page 474
- Adding a receiver VLAN to the MVR VLAN on page 474
- Removing a receiver VLAN from the MVR VLAN on page 475
- Enabling MVR or MVR proxy on a VLAN on page 476
- Enabling MVR or MVR proxy on a VLAN on page 476
- Viewing MVR group information on page 477
- Viewing multicast VLAN information on page 478

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Table 9: Job aid: Roadmap of MVR ACLI commands

Command	Parameter
Global Configuration Mode	
ip igmp multicast-vlan-registration	
ip igmp multicast-vlan-registration	**receiver-vlan <1-4094>**
VLAN Interface Mode	
ip igmp multicast-vlan-registration	**<cr>**
	proxy
Priv EXEC Mode	
show ip igmp multicast-vlan-registration	
show ip igmp multicast-vlan-registration group	**<cr>**
	vlan <1-4094>
	address <A.B.C.D>
ip igmp multicast-vlan-registration vlan	

Enabling MVR globally

Configure MVR to save bandwidth and maintain the performance of the multicast router when multicast groups are located in different VLANs.

Prerequisites

- Disable Protocol Independent Multicast (PIM).
- Log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Enable MVR by using the following command:

   ```
   ip igmp multicast-vlan-registration enable
   ```

2. Display the global MVR configuration by using the following command:

   ```
   show ip igmp multicast-vlan-registration
   ```

Variable definitions

The following table describes optional parameters the you enter after the **`show ip igmp multicast-vlan-registration`** command.

Variable	Value
group	Displays the multicast VLAN group information.
vlan	Displays the multicast VLAN information.
vrf	Displays the multicast VRF information.
vrfids	Displays the multicast VRF ID information.

Disabling MVR globally

Disable MVR so that you can enable PIM.

Prerequisites

- Log on to the Global Configuration mode in the ACLI.

Procedure steps

Disable MVR by using the following command:

```
no ip igmp multicast-vlan-registration enable
```

Configuring MVR on VRF using ACLI

Use this procedure to configure an MVR instance on a VRF.

Prerequisites

- Disable Protocol Independent Multicast (PIM).
- Log on to the VRF Router Configuration mode in ACLI.

Procedure steps

1. Create an MVR instance on a VRF:

   ```
   ip igmp multicast-vlan-registration
   ```

2. Enable the MVR instance on a VRF:

   ```
   ip igmp multicast-vlan-registration enable
   ```

3. To delete the MVR instance on a VRF:

   ```
   no ip igmp multicast-vlan-registration
   ```

4. To disable the non-default MVR instance on a VRF:

   ```
   no ip igmp multicast-vlan-registration enable
   ```

5. To delete a non-default MVR instance on a VRF:

   ```
   default ip igmp multicast-vlan-registration
   ```

6. To enable the default MVR instance on a VRF:

   ```
   default ip igmp multicast-vlan-registration enable
   ```

Adding a receiver VLAN to the MVR VLAN

Add a receiver VLAN to the MVR VLAN to bind the VLAN.

Prerequisites

- Disable PIM.
- Enable MVR globally, when adding the first receiver VLAN. You can subsequently add more receiver VLANs even when MVR is disabled.

- Enable IGMP snoop on a VLAN.
- Log on to the Global Configuration mode in the ACLI.

Procedure steps

Bind a receiver VLAN to the MVR VLAN by using the following command:

```
ip igmp multicast-vlan-registration receiver-vlan <1-4094>
```

Variable definitions

The following table describes variables that you enter in the **`ip igmp multicast-vlan-registration receiver-vlan <1-4094>`** command.

Variable	Value
receiver-vlan <1-4094>	Specifies a VLAN containing a receiver. *<1-4094>* is the identifier for the VLAN. Values range from 1 to 4094.

Removing a receiver VLAN from the MVR VLAN

Remove a receiver VLAN to the MVR VLAN to unbind the VLAN.

Prerequisites

- Disable PIM.
- Log on to the Global Configuration mode in the ACLI.

Procedure steps

Remove a receiver VLAN to the MVR VLAN by using the following command:

```
no ip igmp multicast-vlan-registration receiver-vlan <1-4094>
```

Variable definitions

The following table describes variables that you enter in the **no ip igmp multicast-vlan-registration** command.

Variable	Value
receiver-vlan <1-4094>	Specifies a VLAN containing a receiver.

Enabling MVR or MVR proxy on a VLAN

Enable MVR or MVR Proxy on a VLAN to designate the MVR VLAN for the switch.

Prerequisites

- Disable PIM.
- Enable MVR globally.
- Log on to the Global Configuration mode in the ACLI.

Procedure steps

1. Log on to the VLAN Interface mode in the ACLI by using the following command:

   ```
   interface vlan <1-4094>
   ```

2. Enable a VLAN as the MVR VLAN for the switch by using the following command:

   ```
   ip igmp multicast-vlan-registration
   ```

3. Enable a VLAN as the MVR Proxy VLAN for the switch by using the following command:

   ```
   ip igmp multicast-vlan-registration proxy
   ```

Variable definitions

The following table describes variables that you enter in the **`ip igmp multicast-vlan-registration proxy`** command.

Variable	Value
[proxy]	Enables a VLAN as the MVR Proxy VLAN for the switch.
<1-4094>	Specifies the VLAN ID in the range of 1 to 4094.

Viewing MVR group information

View MVR group information to verify the configuration and troubleshoot.

Prerequisites

- Log on to the Priv EXEC mode in the ACLI.

Procedure steps

View the MVR group by using the following command:

```
show ip igmp multicast-vlan-registration group [vlan <1-4094>]
[address <A.B.C.D>] [vrf <vrf-name>] [vrfids <vrf-ids>]
```

Variable definitions

The following table describes variables that you enter in the **`show ip igmp multicast-vlan-registration group [vlan <1-4094>] [address <A.B.C.D>] [vrf <vrf-name>] [vrfids <vrf-ids>]`** command.

Variable	Value
[vlan <1-4094>]	Specifies the VLAN to which the group belongs.

Variable	Value
	value is the VLAN ID in the range of 1 to 4094.
[address <A.B.C.D>]	Specifies the IP address of a specific group. *value* is the IP address in the format of A.B.C.D.
vrf	Specifies the VRF name.
vrfids	Specifies the VRF IDs.

Viewing multicast VLAN information

View multicast VLAN information to see the VLANs that are bound to the MVR VLAN.

Prerequisites

- Log on to the Priv EXEC mode in the ACLI

Procedure steps

Display the multicast VLAN information by using the following command:

```
show ip igmp multicast-vlan-registration vlan
```

Chapter 29: Multicast MAC filtering using Enterprise Device Manager

With multicast media access control (MAC) filtering, you can create a smaller flooding domain inside a VLAN. For a particular VLAN, you can specify a multicast MAC address and a subset of ports. After clients send data to a designated MAC address, only that subset of ports receives the traffic. For more information about multicast MAC filtering, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Multicast Routing Protocols,*(NN46205-501).

Navigation

- Configuring Layer 2 multicast MAC filtering on page 479
- Configuring Layer 3 multicast MAC filtering on page 480

Configuring Layer 2 multicast MAC filtering

Configure Layer 2 multicast MAC filtering to direct MAC multicast flooding to a specific set of ports.

Prerequisites

- To configure MAC filtering on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. From the table, select a VLAN.
4. Click **Bridge**.

5. Click the **Multicast** tab.
6. Click **Insert**.
7. In the **Address** box, type the MAC address for the multicast flooding domain.
8. Click **[...]** button next to the **ForwardingPorts** box, and then choose from the list of ports that appear.
9. Click **Ok**.
10. Click the ellipsis button **[...]** next to the **MltIds** box, and then choose from the list of MLT IDs that appear.
11. Click **Ok**.
12. Click **Insert**.

Variable definitions

Use the data in the following table to configure the Insert Multicast tab.

Variable	Value

Address	Configures the MAC address for the multicast flooding domain.
	Important:
	The following warning message appears if you attempt to configure the MAC address beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).
	`This MAC address falls in the reserved range for IP multicast. Problems may occur if you have IP multicast configured on this interface.`
	`Do not use any MAC address 01:00:5E:00:00:XX corresponding to the 224.0.0.x reserved IP address range.`

ForwardingPorts	Specifies the ports to include in the multicast flooding domain.
MltIds	Specifies the multilink trunks to include in the multicast flooding domain.

Configuring Layer 3 multicast MAC filtering

Configure Layer 3 multicast MAC filtering to route an IP frame to a unicast IP address and flood it with a destination multicast MAC address. You must manually define a static Address

Resolution Protocol (ARP) entry that associates an IP address with a multicast MAC address, flooding ports, and a multilink trunk.

Prerequisites

- To configure MAC filtering on a specific VRF instance, first change the VRF instance as required.

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Multicast ARP** tab.
4. To add a MAC address, click **Insert**.
5. In the **VlanId** box, click the arrow to choose the VLAN.
6. In the **MacAddress** box, type the MAC address.
7. In the **IpAddress** box, type the IP address.
8. In the **Ports** box, click the [...] button, and then choose from the list of ports that appear.
9. Click **Ok**.
10. Click **Insert**.

Variable definitions

Use the data in the following table to configure the Multicast ARP tab.

Variable	Value
VlanID	Specifies the ID number of the VLAN for the multicast ARP.
MacAddress	Configures the MAC address for the multicast flooding domain. **Important:** The following warning message appears if you attempt to configure the MAC address beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).

Variable	Value
	`This MAC address falls in the reserved range for IP multicast. Problems may occur if you have IP multicast configured on this interface. Do not use any MAC address 01:00:5E:00:00:XX corresponding to the 224.0.0.x reserved IP address range.`
IPAddress	Configures the IP address of the multicast ARP.
Ports	Specifies the ports that receive the multicast flooding.

Chapter 30: Multicast MAC filtering using the CLI

With multicast media access control (MAC) filtering, you can create a smaller flooding domain inside a VLAN. You can specify a multicast MAC address and a subset of ports for a VLAN. After clients send data to that designated MAC address, only that subset of ports receive the traffic.

Multicast MAC filtering navigation

- Job aid on page 483
- Configuring Layer 2 multicast MAC filtering on page 484
- Configuring Layer 3 multicast MAC filtering on page 486

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ip arp static-mcastmac`	`info`
	`add mac <value> ip <value> vlan <value> [port <value>] [mlt <value>]`
	`delete <ipaddr>`
`config vlan <vid> static-mcastmac`	`info`
	`add mac <value> [port <value>] [mlt <value>]`
	`add-mlt <mid> mac <value>`
	`add-ports <ports> mac <value>`
	`delete mac <value>`

Command	Parameter
	delete-mlt <mid> mac <value>
	delete-ports <ports> mac <value>
config ip vrf <vrfName> arp static-mcastmac	**info**
	add mac <value> ip <value> vlan <value> [port <value>] [mlt <value>]
	delete <ipaddr>
config ip vrf <vrfName> vlan <vid> static-mcastmac	**info**
	add mac <value> [port <value>] [mlt <value>]
	add-mlt <mid> mac <value>
	add-ports <ports> mac <value>
	delete mac <value>
	delete-mlt <mid> mac <value>
	delete-ports <ports> mac <value>

Configuring Layer 2 multicast MAC filtering

Configure Layer 2 multicast MAC filtering to direct MAC multicast flooding to a specific set of ports.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **config** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

1. Configure Layer 2 multicast MAC filtering:

   ```
   config vlan <vid> static-mcastmac
   ```

Important:

The following warning message appears if you attempt to configure the **static-mcastmac** parameter with a MAC address beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).

```
This MAC address falls in the reserved range for IP
multicast. Problems may occur if you have IP multicast
configured on this interface.
```

```
Do not use any MAC address 01:00:5E:00:00:XX corresponding
to the 224.0.0.x reserved IP address range.
```

2. Configure the remaining parameters as required.

Variable definitions

The following table describes variables that you enter in the **config vlan <vid> static-mcastmac** command.

Variable	Value
<vid>	Specifies the VLAN ID in the range of 1 to 4094.

Use the data in the following table to use the **config vlan static-mcastmac** command.

Variable	Value
add mac <*value*> [port <*value*>] [mlt <*value*>]	Adds VLAN static multicast MAC entries. • **mac** <*value*> is the MAC address. • **port** <*value*> is the port that receives the multicast flooding. • **mlt** <*value*> is the multilink trunk ID.
add-mlt <*mid*> mac <*value*>	Adds MLT to VLAN static multicast MAC entries. • *mid* is the multilink trunk ID from 1–256. • *value* is the MAC address.
add-ports <*ports*> mac <*value*>	Adds ports to VLAN static multicast MAC entries. • *ports* is a port or a range of ports in slot or port format. • *value* is the MAC address.

Variable	Value
delete mac <*value*>	Deletes VLAN static multicast MAC entries. • *value* is the MAC address.
delete-mlt <*mid*> mac <*value*>	Deletes MLT-to-VLAN static multicast MAC entries. • *mid* is the multilink trunk ID from 1–256. • **mac** <*value*> is the MAC address.
delete-ports <*ports*> mac <*value*>	Deletes ports from VLAN static multicast MAC entries. • *ports* is a port or a range of ports in slot or port format. • *value* is the MAC address.
info	Displays current settings.

Example of configuring Layer 2 multicast MAC filtering

Procedure steps

Add a multicast MAC address 01:02:03:04:05:06 as a static MAC in VLAN 2. Add ports and a multilink trunk group so that traffic destined for the MAC address is forwarded to ports 4/1 through 4/4 and MLT 1, instead of flooding to all VLAN 2 ports.

```
ERS-8606:5# config vlan 3 static-mcastmac add mac 01:02:03:04:05:06 port
4/1-4/4 mlt 1
```

Configuring Layer 3 multicast MAC filtering

Configure Layer 3 multicast MAC filtering to route an IP frame to a unicast IP address and flood it with a destination multicast MAC address. You must manually define a static Address Resolution Protocol (ARP) entry that associates an IP address with a multicast MAC address, flooding ports, and a multilink trunk.

You use mroute commands on a VRF the same way you configure for the Global Router, except that you must replace **config ip** with **config ip vrf <vrfName>** in the following procedure.

Procedure steps

Configure Layer 3 multicast MAC filtering:

```
config ip arp static-mcastmac
```

Important:

The following warning message appears if you attempt to configure the **static-mcastmac** parameter with a MAC address beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).

```
This MAC address falls in the reserved range for IP
multicast. Problems may occur if you have IP multicast
configured on this interface.
```

```
Do not use any MAC address 01:00:5E:00:00:XX corresponding to
the 224.0.0.x reserved IP address range.
```

Variable definitions

Use the data in the following table to use the **config ip arp static-mcastmac** and **config ip vrf <vrfName> arp static-mcastmac** commands.

Variable	Value
add mac <*value*> ip <*value*> vlan <*value*> [port <*value*>] [mlt <*value*>]	Adds static multicast MAC entries. • **mac** <*value*> is the MAC address. • **ip** <*value*> is the IP address. • **vlan** <*value*> is the VLAN ID number. • **port** <*value*> is the port that receives the multicast flooding. • **mlt** <*value*> is the multilink trunk ID.
delete <*ipaddr*>	Deletes static multicast MAC entries. • *ipaddr* is the IP address.
info	Displays current settings.

Example of configuring Layer 3 multicast MAC filtering

Procedure steps

Add a multicast MAC address 01:01:01:01:01:02 as a static ARP entry in VLAN 2. Add ports and a multilink trunk group so that traffic destined for the MAC address is forwarded to ports 4/14 and 4/43, and MLT 1, instead of flooding to all VLAN 2 ports.

```
ERS-8606:5# config ip arp static-mcastmac add mac 01:01:01:01:01:02 ip
2.2.2.100 vlan 2 port 4/14-4/43 mlt 1
```

Chapter 31: Multicast MAC filtering using the ACLI

With multicast media access control (MAC) filtering, you can create a smaller flooding domain inside a VLAN. You can specify a multicast MAC address and a subset of ports for a VLAN. After clients send data to that designated MAC address, only that subset of ports receive the traffic.

Multicast MAC filtering navigation

- Job aid on page 489
- Configuring Layer 2 multicast MAC filtering on page 489
- Configuring Layer 3 multicast MAC filtering on page 491

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command
Global Configuration and VRF Router Configuration modes
ip arp static-mcast <A.B.C.D> <MAC address> vid <1-4094> [port <value> Word<1-16>]
vlan mac-address-static <1-4094> <MAC address> <ports> qos <0-7>

Configuring Layer 2 multicast MAC filtering

Configure Layer 2 multicast MAC filtering to direct MAC multicast flooding to a specific set of ports.

You configure MAC filtering on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Configure Layer 2 multicast MAC filtering:

```
vlan mac-address-static <1-4094> <MAC address> <ports> qos
<0-7>
```

Important:

The following warning message appears if you attempt to configure the **`mac-address-static`** parameter with a MAC address beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).

```
This MAC address falls in the reserved range for IP
multicast. Problems may occur if you have IP multicast
configured on this interface.
```

```
Do not use any MAC address 01:00:5E:00:00:XX corresponding to
the 224.0.0.x reserved IP address range.
```

Variable definitions

Use the data in the following table to use the **`vlan mac-address-static`** command.

Variable	Value
mac address	Specifies the MAC address in hexadecimal format.
<ports>	Specifies the port or ports that receive the multicast flooding. Type this information as a port or a range of ports in slot or port format.
qos <0-7>	Specifies the Quality of Service level.
1–4094	Specifies a VLAN from 1–4094.

Example of configuring Layer 2 multicast MAC filtering

Procedure steps

Add a multicast MAC address 01:02:03:04:05:06 as a static MAC in VLAN 2. Add ports so that traffic destined for the MAC address is forwarded to ports 4/1 through 4/4, instead of flooding to all VLAN 2 ports.

```
ERS-8606:5(config)#vlan mac-address-static 2 01:02:03:04:05:06 4/1-4/4 qos 2
```

Configuring Layer 3 multicast MAC filtering

Configure Layer 3 multicast MAC filtering to route an IP frame to a unicast IP address and flood it with a destination multicast MAC address. You must manually define a static ARP entry that associates an IP address with a multicast MAC address, flooding ports, and a multilink trunk.

You configure MAC filtering on a VRF instance the same way you configure for the Global Router, except that you must use VRF Router Configuration mode.

Prerequisites

- You must log on to the Global Configuration mode in the ACLI.

Procedure steps

Configure Layer 3 multicast MAC filtering:

```
ip arp static-mcast <A.B.C.D> <MAC address> vid <1-4094> [port
<value> Word<1-16>]
```

Important:

The following warning message appears if you attempt to configure the **static-mcast** parameter with a MAC address beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).

```
This MAC address falls in the reserved range for IP
multicast. Problems may occur if you have IP multicast
configured on this interface.
```

```
Do not use any MAC address 01:00:5E:00:00:XX corresponding to
the 224.0.0.x reserved IP address range.
```

Variable definitions

Use the data in the following table to use the **ip arp static-mcast <A.B.C.D> <MAC address> vid <1–4094> [port <value> Word<1–16>]** command.

Variable	Value
<mac address>	Specifies the MAC address in hexadecimal format.
<A.B.C.D>	Specifies the IP address.
vid <1–4094>	Specifies the VLAN ID.
port <value>	Specifies the port that receives the multicast flooding. Type this information as a port or a range of ports in slot or port format.
<1–16>	Specifies the multilink trunk ID.

Example of configuring Layer 3 multicast MAC filtering

Procedure steps

Add a multicast MAC address 01:01:01:01:01:02 as a static ARP entry in VLAN 2. Add ports and a multilink trunk group so that traffic destined for the MAC address is forwarded to ports 4/14 and 4/43, and MLT 1, instead of flooding to all VLAN 2 ports.

```
ERS-8606:5(config)# ip arp static-mcast 01:01:01:01:01:02 2.2.2.100 vid 2
port 4/14-4/43 1
```

Chapter 32: Common procedures using Enterprise Device Manager

The following sections describe common procedures that you use while configuring multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

Navigation

- Configuring a prefix list on page 493

Configuring a prefix list

Configure a prefix list to use in multicast access control and routing policies. Prefix lists are lists of routes that apply to one or more route policies. They contain a set of contiguous or noncontiguous routes. Reference prefix lists by name from within the routing policies.

For more information about prefix lists, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — IP Routing,* (NN46205-523).

Procedure steps

1. In the navigation tree, open the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click **Insert**.
4. Edit the appropriate data.
5. Click **Insert**.

Variable definitions

Use the data in the following table to configure the Policy, Insert Prefix List dialog box.

Field	Description
ID	The list identifier.
Prefix	The IP address.
PrefixMaskLen	The length of the prefix mask. **Important:** You must use the full 32-bit mask to exact a full match of a specific IP address (for example, when you create a policy to match on next hop).
Name	The name of a specified prefix list during the creation process, or renames the specified prefix list. The name can contain 1–64 characters.
MaskLenFrom	The lower boundary of the mask length. The default is the mask length. Lower boundary and higher boundary mask lengths together can define a range of networks.
MaskLenUpto	The higher boundary of the mask length. The default is the mask length. Lower boundary and higher boundary mask lengths together can define a range of networks.

Chapter 33: Common procedures using the CLI

The following sections describe common procedures that you use while configuring multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

- Job aid on page 495
- Creating an IP prefix list on page 495

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
`config ip prefix-list <prefix-list-name>`	`info`
	`add-prefix <ipaddr|mask> [maskLenFrom <value>] [maskLenTo <value>]`
	`delete`
	`name <name>`
	`remove-prefix <ipaddr/mask>`

Creating an IP prefix list

You can create one or more IP prefix lists and apply those lists to an IP route policy. A prefix list with a 32-bit mask is equivalent to an address. You can use a prefix list with a mask less than 32 bits as a network. If you configure the MaskLenFrom field lower than the MaskLenUpto field, you can also use it as a range.

Procedure steps

Create a prefix list:

```
config ip prefix-list <prefix-list-name>
```

Variable definitions

Use the data in the following table to use the **`config ip prefix-list`** command.

Variable	Value
add-prefix <*ipaddr*|*mask*> [maskLenFrom <*value*>] [maskLenTo <*value*>]	Adds a prefix entry to the prefix list. • *ipaddr*|*mask* is the IP address and mask. • **`maskLenFrom`** <*value*> is the lower boundary of the mask length. The default is the mask length. • **`maskLenTo`** <*value*> is the higher boundary of the mask length. The default is the mask length. **Important:** Lower boundaries and higher boundaries of the mask lengths together can define a range of networks.
delete	Deletes the prefix list.
info	Displays all of the prefixes in a list.
name <*name*>	Renames the specified prefix list. The name can contain 1–64 characters.
prefix-list-name	Indicates the name of the specified prefix list, which is a string length from 1–64 characters.
remove-prefix <*ipaddr/mask*>	Removes a prefix entry from the prefix list. • *ipaddr/mask* is the IP address and mask.

Example of creating an IP prefix list

Procedure steps

1. Create a prefix list.

```
ERS_8606:5#config vlan 3 ip config ip prefix-list prefix1
```

2. Add a prefix entry to the list.

```
ERS_8606:5/config/ip/prefix-list/prefix1# add 4.4.4.4/24
```

Chapter 34: Common procedures using the ACLI

The following sections describe common procedures that you use while configuring multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

- Job aid on page 499
- Creating an IP prefix list on page 499

Job aid

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Global Configuration mode	
ip prefix-list <prefix-list-name>	**name <name>**
	IP address

Creating an IP prefix list

You can create one or more IP prefix lists and apply those lists to an IP route policy. A prefix list with a 32-bit mask is equivalent to an address. You can configure a prefix list with a mask less than 32 bits as a network.

Prerequisites

- You must log on to the Global Configuration mode.

Procedure steps

Create a prefix list by naming it and adding a prefix entry:

```
ip prefix-list <prefix-list-name> <IP address>
```

Variable definitions

Use the data in the following table to use the **ip prefix-list** command.

Variable	Value
name *<name>*	Renames the specified prefix list. The name can contain 1–64 characters.
prefix-list-name	Indicates the name of the specified prefix list, which is a string length from 1–64 characters. Use the **no** operator to later remove this configuration.
IP address	Adds a prefix entry to the prefix list. Specify the IP address in the format a.b.c.d/x or a.b.c.d/x.x.x.x **Important:** Lower boundaries and higher boundaries of the mask lengths together can define a range of networks. Use the **no** operator to later remove this configuration.

Example of creating an IP prefix list

Procedure steps

Create a prefix list and add a prefix entry to the list.

```
ERS-8606:5(config)#ip prefix-list prefix1 4.4.4.4/24
```

Chapter 35: CLI show command reference

This reference information provides show commands to view the operational status of multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

CLI show command reference navigation

- DVMRP on page 502
- DVMRP interface on page 503
- DVMRP neighbor on page 504
- DVMRP ports on page 505
- DVMRP route on page 506
- DVMRP VLAN on page 507
- IGMP access on page 508
- IGMP cache on page 509
- IGMP group on page 509
- IGMP interface on page 510
- IGMP multicast router discovery on page 511
- IGMP multicast router discovery neighbors on page 512
- IGMP ports on page 512
- IGMP router-alert on page 514
- IGMP sender on page 514
- IGMP snoop on page 515
- IGMP static and blocked ports on page 515
- IGMP VLAN on page 516
- Layer 2 multicast MAC filters on page 517
- Layer 3 multicast MAC ARP data on page 518
- Multicast group trace for IGMP snoop on page 518
- Multicast MLT distribution on page 519

- Multicast route information on page 520
- Multicast route next hop on page 522
- Multicast routes on an interface on page 523
- Multicast static route information on page 523
- Multicast VLAN information on page 524
- PGM global information on page 535
- PIM active RP on page 536
- PIM bootstrap router on page 537
- PIM interface on page 538
- PIM mode on page 539
- PIM neighbor on page 540
- PIM route on page 540
- PIM virtual neighbor on page 542
- Rendezvous points on page 542
- SSM channel information on page 543
- SSM group range and dynamic learning status on page 544
- Static RP table on page 544
- Static source groups on page 545
- VLAN port data on page 546

DVMRP

Use the **`show ip dvmrp info`** command to display information about the general Distance Vector Multicast Routing Protocol (DVMRP) group. The syntax for this command is as follows.

```
show ip dvmrp info
```

The following table shows the field descriptions for this command.

Table 10: show ip dvmrp info command

Field	Description
AdminStat	Indicates the status of DVMRP.

Field	Description
Genid	Indicates the generation identifier for the routing process. This ID is used by neighboring routers to detect whether the DVMRP routing table is present.
Version	Indicates the version of DVMRP.
NumRoutes	Indicates the number of entries in the routing table. This value is used to monitor the routing table size to detect illegal advertisements of unicast routes.
NumReachableRoutes	Indicates the number of entries in the routing table with noninfinite metrics. This value is used to detect network partitions by observing the ratio of reachable routes to total routes.
UpdateInterval	Indicates the global route update interval in seconds.
TriggeredUpdateInterval	Indicates the global route triggered update interval in seconds.
LeafTimeOut	Indicates the hold down timer for leaf in seconds.
NbrTimeOut	Indicates the neighbor timeout interval in seconds.
NbrProbeInterval	Indicates the global neighbor probe interval in seconds.
FwdCacheTimeout	Indicates the timeout interval (in seconds) for aging prune entries in the forward cache.
RouteExpireTimeout	Indicates the minimum amount of time remaining before this entry ages out.
RouteDiscardTimeout	Indicates the interval (in seconds) to discard collected routes.
RouteSwitchTimeout	Indicates the interval (in seconds) to discard unused routes.
ShowNextHopTable	Indicates the status of showing the next-hop table.
generate-trap	Indicates the status of DVMRP traps.
generate-log	Indicates the status of logging DVMRP.
PruneResend	Indicates the status of resending prune messages.

DVMRP interface

Use the **`show ip dvmrp interface`** command to display DVMRP route policy information for the DVMRP interface configurations on the switch. The syntax for this command is as follows.

```
show ip dvmrp interface
```

The following table shows the field descriptions for this command.

Table 11: show ip dvmrp interface command

Field	Description
IF	Indicates the ifIndex value of the interface for which DVMRP is enabled.
ADDR	Indicates the IP address this system uses as a source address on this interface.
METRIC	Indicates the distance metric for this interface used to calculate distance vectors.
OPERSTAT	Indicates the current operational state of this DVMRP interface.
DEFAULT LISTEN	Indicates whether the switch can learn DVMRP default routes over this interface.
DEFAULT SUPPLY	Indicates if the switch supplies DVMRP default routes over this interface.
DEFAULT METRIC	Indicates the cost of the DVMRP default route that this interface generates and supplies when configured to supply the default route.
ADVERTISE SELF	Indicates whether the switch can advertise this local network.
IN-POLICY	Indicates the DVMRP accept policy name configured on this interface.
OUT-POLICY	Indicates the DVMRP announce policy name configured on this interface.
INTF TYPE	Indicates the type of this DVMRP interface and whether it uses a tunnel, source routing, a physical interface for which a querier exists, or a physical interface for which a querier does not exist (subnet).

DVMRP neighbor

Use the **`show ip dvmrp neighbor`** command to display information about the configured DVMRP neighbors. The syntax for this command is as follows.

```
show ip dvmrp neighbor
```

The following table shows the field descriptions for this command.

Table 12: show ip dvmrp neighbor command

Field	Description
INTERFACE	Indicates the value of ifIndex for the virtual interface used to reach this DVMRP neighbor.
ADDRESS	Indicates the IP address of the DVMRP neighbor.
EXPIRE	Indicates the minimum time (in seconds) remaining before this DVMRP neighbor ages out.
GENID	Indicates the neighboring router generation identifier.
MAJVER	Indicates the major DVMRP version number of the neighboring router.
MINVER	Indicates the minor DVMRP version number of the neighboring router.
CAPABILITY	Indicates the capabilities of the neighboring router. The leaf bit indicates that the neighbor uses only one interface with its neighbors. The prune bit indicates that the neighbor supports pruning. The generation ID bit indicates that the neighbor sends its generation ID in probe messages. The mtrace bit indicates that the neighbor can handle mtrace requests.
STATE	Indicates the state of DVMRP for the neighboring router.

DVMRP ports

Use the **`show ports info dvmrp`** command to display DVMRP route policy information for the DVMRP port configurations on the switch. The syntax for this command is as follows.

```
show ports info dvmrp [port <value>]
```

Use the data in the following table to use the **`show ports info dvmrp`** command.

Table 13: Command parameters

Variable	Value
`port <value>`	Specifies a port or range of ports in the slot or port format.

The following table shows field descriptions for this command.

Table 14: show ports info dvmrp command

Field	Description
PORT NUM	Indicates the slot and port number associated with the interface.
DVMRP ENABLE	Indicates if DVMRP is enabled or disabled on the port.
METRIC	Indicates the distance metric for this port used to calculate distance vectors.
DEFAULT LISTEN	Indicates whether the switch can learn DVMRP default routes over this port.
DEFAULT SUPPLY	Indicates if the switch supplies DVMRP default routes over this port.
DEFAULT METRIC	Indicates the cost of the DVMRP default route that this port generates and supplies when configured to supply the default route.
ADVERTISE SELF	Indicates whether the switch can advertise this local network.

DVMRP route

Use the **`show ip dvmrp route`** command to display information about the DVMRP routes. The syntax for this command is as follows.

```
show ip dvmrp route
```

The following table shows the field descriptions for this command.

Table 15: show ip dvmrp route command

Field	Description
SOURCE	Indicates the network address that when combined with the corresponding value of dvmrpRouteSourceMask, identifies the sources for multicast routing information.
MASK	Indicates the network mask that when combined with the corresponding value of dvmrpRouteSource, identifies the sources for multicast routing information.
UPSTREAM_NBR	Indicates the address of the upstream neighbor (for example, RPF neighbor) from which IP datagrams from these sources are received.

Field	Description
INTERFACE	Indicates the interface on which IP datagrams sent by these sources are received.
METRIC	Indicates the distance in hops to the source subnet.
EXPIRE	Indicates the minimum amount of time remaining (in seconds) before this entry ages out.

DVMRP VLAN

Use the **show vlan info dvmrp** command to display DVMRP route policy information for the DVMRP VLAN configurations on the switch. The syntax for this command is as follows.

```
show vlan info dvmrp [vlan <value>]
```

Use the data in the following table to use the **show vlan info dvmrp** command.

Table 16: Command parameters

Parameters	Value
`vlan <value>`	Specifies a VLAN ID from 1–4092.

See the following table for field descriptions for this command.

Table 17: show vlan info dvmrp command

Field	Description
VLAN ID	Indicates the VLAN ID associated with the VLAN interface.
DVMRP ENABLE	Indicates if DVMRP is enabled or disabled on the VLAN interface.
METRIC	Indicates the distance metric for this VLAN interface used to calculate distance vectors.
DEFAULT LISTEN	Indicates whether the switch can learn DVMRP default routes over this VLAN interface.
DEFAULT SUPPLY	Indicates if the switch supplies DVMRP default routes over this VLAN interface.
DEFAULT METRIC	Indicates the cost of the DVMRP default route that this VLAN interface generates and supplies when configured to supply the default route.

Field	Description
ADVERTISE SELF	Indicates whether the switch can advertise this local network.

IGMP access

Use the **show ip igmp access** command to display information about the Internet Group Management Protocol (IGMP) multicast access control groups. The syntax for this command is as follows.

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

```
show ip igmp access [vrf <value>] [vrfids <value>]
```

The following table shows the field descriptions for this command.

Table 18: show ip igmp access field descriptions

Field	Description
INTERFACE	Identifies the interface where multicast access control is configured.
GRP PREFIX	An alphanumeric string that identifies the name of the access policy.
HOSTADDR	The IP address of the host.
HOSTMASK	The subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the host network.
ACCESSMODE	Specifies the action of the access policy. The actions are • deny-tx—deny IP multicast transmitted traffic • deny-rx—deny IP multicast received traffic • deny-both—deny both IP multicast transmitted and received traffic • allow-only-rx—allow IP multicast transmitted traffic • allow-only-rx—allow IP multicast received traffic • allow-only-both—allow both IP multicast transmitted and received traffic

IGMP cache

Use the **show ip igmp cache** command to display information about the IGMP cache. The syntax for this command is as follows.

```
show ip igmp cache [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 19: show ip igmp cache command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INTERFACE	Indicates the physical interface or a logical interface (VLAN) that received group reports from various sources.
LASTREPORTER	Indicates the IP address of the source of the last membership report received for this IP multicast group address on this interface. If no membership report is received, this object uses the value 0.0.0.0.
EXPIRATION	Indicates the minimum amount of time remaining before this entry ages out.
V1HOSTTIMER	Indicates the time remaining until the local router assumes that no IGMPv1 members exist on the IP subnet attached to this interface.
TYPE	Indicates whether the entry is learned dynamically or is added statically.
STATICPORTS	Indicates the IP multicast group address for which this entry contains information.

IGMP group

Use the **show ip igmp group** command to display information about the IGMP group. The syntax for this command is as follows.

```
show ip igmp group [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 20: show ip igmp group command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INPORT	Indicates the physical interface or a logical interface (VLAN) that received group reports from various sources.
MEMBER	Indicates the IP address of a source that sent a group report to join this group.
EXPIRATION TIME	Indicates the time left before the group report expires on this port. This variable is updated upon receiving a group report.

IGMP interface

Use the **show ip igmp interface** command to display information about the interfaces where IGMP is enabled. This syntax for this command is as follows.

```
show ip igmp interface [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 21: show ip igmp interface command

Field	Description
IF	Indicates the interface where IGMP is configured.
QUERY INTVL	Indicates the frequency at which the interface transmits IGMP host query packets.
STATUS	Indicates the activation of a row that enables IGMP on the interface. The destruction of a row disables IGMP on the interface.

Field	Description
VERS.	Indicates the version of IGMP that is running on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface is attached.
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received whose IGMP version does not match the igmpInterface version. You must configure all routers on a LAN to run the same version of IGMP. Thus, if the interface receives queries with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times IGMP added a group membership on this interface.
ROBUST	Indicates the robustness variable, which you configure for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.
LASTMEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This does not apply if igmpInterface version is 1.

IGMP multicast router discovery

Use the **`show ip igmp mrdisc`** command to display information about the IGMP multicast discovery routes. The syntax for this command is as follows.

```
show ip igmp mrdisc [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 22: show ip igmp mrdisc command

Field	Description
VLAN ID	Indicates the VLAN ID.
MRDISC	Indicates the status of multicast router discovery.
DISCOVERED RTR PORTS	Indicates the ports discovered.

IGMP multicast router discovery neighbors

Use the **`show ip igmp mrdisc-neighbors`** command to display information about the IGMP multicast router discovery neighbors. The syntax for this command is as follows.

```
show ip igmp mrdisc-neighbors [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 23: show ip igmp mrdisc-neighbors command

Field	Description
VLAN ID	Indicates the VLAN ID.
SRC_PORT	Indicates the source port.
IP Addr	Indicates the IP address.
Advert-int	Indicates the advertisement interval in seconds.
QUERY-int	Indicates the query interval in seconds.
Robust-val	Indicates the tuning for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the Robustness variable.

IGMP ports

Use the **`show ports info igmp`** command to display information about the specified port or for all ports. The syntax for this command is as follows.

```
show ports info igmp [<ports>] [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Use the data in the following table to use the **show ports info igmp** command.

Table 24: Command parameters

Parameter	Value
ports	Identifies a specific port using the convention {slot/port[-slot/port][, ...]}.

The following table shows the field descriptions for this command.

Table 25: show ports info igmp command

Field	Description
PORT NUM	Indicates the port number.
QUERY INTVL	Indicates the interval (in seconds) between IGMP host query packets transmitted on this port.
QUERY MAX RESP	Indicates the interval (in seconds) for the maximum query response time advertised in IGMPv2 queries on this interface. Smaller values allow a router to prune groups faster.
ROBUST	Indicates the tuning for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that is running on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response in a group-specific query.
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop on the port.
SNOOP ENABLE	Indicates the status of IGMP snoop on the port.
SSM SNOOP ENABLE	Indicates the status of SSM IGMP snoop on the port.
FAST LEAVE ENABLE	Indicates the status of fast leave.

IGMP router-alert

Use the **show ip igmp router-alert** command to display the status of IGMP router alert. The syntax for this command is as follows.

```
show ip igmp router-alert [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 26: show ip igmp router-alert command

Field	Description
IFINDEX	Indicates the interface index number.
ROUTER ALERT ENABLE	Indicates the status of the router alert check.

IGMP sender

Use the **show ip igmp sender** command to display information about the IGMP senders. The syntax for this command is as follows.

```
show ip igmp sender [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 27: show ip igmp sender command

Field	Description
GRPADDR	Indicates the IP multicast address.
IFINDEX	Indicates the interface index number.
MEMBER	Indicates the IP address of the host.
PORT/MLT	Indicates the IGMP sender ports.

Field	Description
STATE	Indicates if a sender exists because of an IGMP access filter. Options include filtered and nonfiltered.

IGMP snoop

Use the **show ip igmp snoop** command to display the status of IGMP snoop. The syntax of this command is as follows.

```
show ip igmp snoop [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 28: show ip igmp snoop command

Field	Description
IFINDEX	Indicates the interface index number.
SNOOP ENABLE	Indicates the status of IGMP snoop.
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop.
SSM SNOOP ENABLE	Indicates the status of IGMP Source Specific Multicast (SSM) snoop.
STATIC MROUTER PORTS	Indicates the set of ports in this VLAN that provide connectivity to an IP multicast router.
ACTIVE MROUTER PORTS	Indicates the active ports.
MROUTER EXPIRATION TIME	Indicates the multicast querier router aging timeout in seconds.

IGMP static and blocked ports

Use the **show ip igmp static** command to display information about the static and blocked ports for the IGMP-enabled interfaces. The syntax for this command is as follows.

```
show ip igmp static [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 29: show ip igmp static command

Field	Description
GRPADDR	Indicates the IP multicast address. The group address holds the starting range for the address range.
TO-GRPADDR	Indicates the end of the range for the group address.
INTERFACE	Indicates the interface IP address.
STATICPORTS	Indicates the egressing ports.
BLOCKEDPORTS	Indicates the ports not allowed to join.

IGMP VLAN

Use the **`show vlan info igmp`** command to display the IGMP configuration information for all VLANs on the switch or for a specified VLAN. The syntax for this command is as follows.

```
show vlan info igmp [<vid>] [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Use the data in the following table to use the **`show vlan info igmp`** command.

Table 30: Command parameters

Parameter	Value
`vid`	Specifies a VLAN ID from 1–4092.

The following table shows the field descriptions for this command.

Table 31: show vlan info igmp command

Field	Description
VLAN ID	Indicates the VLAN ID.
QUERY INTVL	Indicates the interval (in seconds) between IGMP host query packets transmitted on this interface.

Field	Description
QUERY MAX RESP	Indicates the interval (in seconds) for the maximum query response time advertised in IGMPv2 queries on this interface. Smaller values allow a router to prune groups faster.
ROBUST	Indicates the tuning for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that is running on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the max response in a group specific query.
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop on the VLAN.
SNOOP ENABLE	Indicates the status of IGMP snoop on the VLAN.
SSM SNOOP ENABLE	Indicates the status of SSM IGMP snoop on the VLAN.
FAST LEAVE ENABLE	Indicates the status of fast leave.
FAST LEAVE PORTS	Indicates the ports with fast leave enabled.

Layer 2 multicast MAC filters

Use the **`show vlan info static-mcastmac`** command to display the Layer 2 multicast MAC filters. The syntax for this command is as follows.

```
show vlan info static-mcastmac [<vid>]
```

Use the data in the following table to use the **`show vlan info static-mcastmac`** command.

Table 32: Command parameters

Parameter	Value
`vid`	Specifies the VLAN ID from 1–4092. The VLAN ID is optional. If you use a VLAN ID, the command displays information for the specified VLAN. Without the VLAN ID, the command displays information for all the configured VLANs.

The following table shows the field descriptions for this command.

Table 33: show vlan info static-mcastmac command

Field	Description
VLAN ID	Indicates the VLAN ID.
MAC ADDRESS	Indicates the MAC address.
PORT LIST	Indicates the list of ports.
MLT GROUPS	Indicates the multilink trunk groups.

Layer 3 multicast MAC ARP data

Use the **show ip arp static-mcastmac** command to display Layer 3 multicast MAC ARP data. The syntax for this command is as follows.

```
show ip arp static-mcastmac [<ip address>] [-s <value>] [vrf <value>]
[vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 34: show ip arp static-mcastmac command

Field	Description
IP_ADDRESS	Indicates the multicast IP address.
MAC ADDRESS	Indicates the multicast MAC address.
VLAN	Indicates the VLAN ID.
PORT LIST	Indicates the list of ports.
MLT ID	Indicates the multilink trunk ID.

Multicast group trace for IGMP snoop

Use the **show ip igmp snoop-trace** command to view multicast group trace information for IGMP snoop. Multicast group trace tracks the data flow path of the multicast streams. This command provides information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port. The syntax for the command is as follows.

```
show ip igmp snoop-trace [src <value>] [grp <value>] [vrf <value>]
[vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table describes the command parameters.

Table 35: Command parameters

Parameter	Description
`grp <value>`	Specifies the source IP address in the format a.b.c.d.
`src <value>`	Specifies the group IP address in the format a.b.c.d.

The following table provides the field descriptions for this command.

Table 36: show ip igmp snoop-trace command

Field	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.
OUT PORT	Indicates the outgoing port number.

Multicast MLT distribution

Use the **`show sys mcast-mlt-distribution`** command to show the multicast MultiLink Trunking (MLT) distribution configuration. Use the **`show mlt info`** command to show multicast flow distribution for each multilink trunk. The syntax for these commands is as follows.

```
show sys mcast-mlt-distribution
```

```
show mlt info
```

The following table shows the field descriptions for this command.

Table 37: show mlt info command

Field	Description
MLTID	Indicates the multilink trunk ID number.
IFINDEX	Indicates the interface index number.
NAME	Indicates the name of this multilink trunk.
PORT TYPE	Indicates the type of multilink trunk port: access or trunk.
SVLAN TYPE	Indicates the type of multilink trunk port.
MLT ADMIN	Indicates the status of MLT.
MLT CURRENT	Indicates the operational status of MLT.
PORT MEMBERS	Indicates the set of ports that are members of this multilink trunk.
VLAN IDS	Indicates the number of VLANs on the multilink trunk.
MULTICAST DISTRIBUTION	Indicates the status of multicast distribution for each multilink trunk.
NT-STG	Indicates whether this spanning tree group (STG) is operating in Avaya mode or in Cisco mode. • true—Avaya mode • false—Cisco mode
DESIGNATED PORTS	Indicates the designated port for the multilink trunk.
LACP ADMIN	Indicates the administrative status of link aggregation on the multilink trunk.
LACP OPER	Indicates the operational status of link aggregation on the multilink trunk.

Multicast route information

Use the **`show ip mroute route`** command to display information about the multicast routes set up on the switch. The syntax for this command is as follows.

```
show ip mroute route [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following figure shows sample output for the **`show ip mroute route`** command. In this table, every stream uses one (*,G) entry and x (S,G) entries, depending on how many servers forward traffic to the same group. The 0.0.0.0 mask is always tied to a (*,G) entry. Every time

a new stream comes in, Protocol Independent Multicast (PIM) creates two entries in the table; one is a (*,G) entry that points toward the rendezvous point router, and the other is an (S,G) entry that points toward the source.

```
==========================================================================
                              Mroute Route
==========================================================================
GROUP          SOURCE          SRCMASK        UPSTREAM_NBR    IF    EXPIR PROT

224.2.128.119 0.0.0.0         0.0.0.0        104.0.0.207     v504 210   pimsm
224.2.128.119 155.120.50.165 255.255.255.0 104.0.0.207     v504 210   pimsm
224.2.128.119 155.120.51.165 255.255.255.0 155.120.51.204 v155 210   pimsm
224.2.128.119 155.120.51.166 255.255.255.0 155.120.51.204 v155 210   pimsm
224.2.148.113 0.0.0.0         0.0.0.0        104.0.0.208     v504 210   pimsm
224.2.148.113 155.120.50.165 255.255.255.0 104.0.0.208     v504 210   pimsm
224.2.157.11  0.0.0.0         0.0.0.0        104.0.0.208     v504 210   pimsm
224.2.157.11  155.120.50.165 255.255.255.0 104.0.0.207     v504 210   pimsm
224.2.157.11  155.120.51.165 255.255.255.0 155.120.51.204 v155 210   pimsm
224.2.157.11  155.120.51.166 255.255.255.0 155.120.51.204 v155 210   pimsm
```

Figure 25: show ip mroute route output

The following table shows the field descriptions for this command.

Table 38: show ip mroute route command

Field	Description
GROUP	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
SOURCE	Indicates the network address that, when combined with the corresponding value of ipMRouteNextHopSourceMask, identifies the sources for which this entry specifies a next hop on an outgoing interface.
SRCMASK	Indicates the network mask that, when combined with the corresponding value of ipMRouteNextHopSource, identifies the sources for which this entry specifies a next hop on an outgoing interface.
UPSTREAM_NBR	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is unknown.
IF	Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).
EXPIR	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
PROT	Indicates the outgoing mechanism through which the switch learns this route.

Multicast route next hop

Use the **show ip mroute next-hop** command to display information about the next hop for the multicast routes set up on the switch. The syntax for this command is as follows.

```
show ip mroute next-hop [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 39: show ip mroute next-hop command

Field	Description
INTERFACE	Indicates the interface identity.
GROUP	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
SOURCE	Indicates the network address, which when combined with the corresponding value of ipMRouteNextHopSourceMask identifies the sources for which this entry specifies a next hop on an outgoing interface.
SRCMASK	Indicates the network mask, which when combined with the corresponding value of ipMRouteNextHopSource identifies the sources for which this entry specifies a next hop on an outgoing interface.
ADDRESS	Indicates the address of the next hop specific to this entry. For most interfaces, this address is identical to ipMRouteNextHopGroup. Non Broadcast Multiple Access (NBMA) interfaces, however, use multiple next-hop addresses on a single outgoing interface.
STATE	Indicates whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. The value forwarding indicates the information is currently used; the value pruned indicates it is not used.
EXPTIME	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
CLOSEHOP	Indicates the minimum number of hops between this router and members of this IP multicast group reached through this next hop on this outgoing interface. IP multicast datagrams for the group that use a TTL less than this number of hops are forwarded to the next hop
PROTOCOL	Indicates the routing mechanism through which the switch learns this next hop.

Multicast routes on an interface

Use the **show ip mroute interface** command to display information about the multicast routes set up on the switch for a specific interface. The syntax for this command is as follows.

```
show ip mroute interface [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 40: show ip mroute interface command

Field	Description
INTERFACE	Indicates the interface.
TTL	Indicates the datagram TTL threshold for the interface. The interface does not forward IP multicast datagrams with a TTL less than this threshold. The default value of 0 means the interface forwards all multicast packets.
PROTOCOL	Indicates the routing protocol running on this interface.

Multicast static route information

Use the **show ip mroute rpf** command to view the best route to reach a specific RPF address from either the static multicast or unicast routing table. The syntax for this command is as follows.

```
show ip mroute rpf <addr> [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Important:

Because the switch uses two static IP route tables, one for unicast routes and one for multicast routes, the command **show ip route info** does not provide the correct RPF information to a source or RP. Avaya recommends that you use the **show ip mroute rpf** command.

Multicast VLAN information

Use the **show vlan info all** command to display information about all the VLANs on the switch. The syntax for this command is as follows.

```
show vlan info all [<vid>] [port <value>] [by <value>]
```

Use the data in the following table to use the **show vlan info all** command.

Table 41: Command parameters

Parameter	Value
by <value>	Specifies the group ID.
port <value>	Specifies the port or range of ports in slot or port format.
vid	Specifies the VLAN ID from 1–4092. The VLAN ID is optional. If you use a VLAN ID, the command displays information for the specified VLAN. Without the VLAN ID, the command displays information for all the configured VLANs.

The following table shows the field descriptions for this command.

Table 42: show vlan info all command

Field	Description
Vlan Basic	
VLAN ID	Indicates the VLAN ID.
NAME	Indicates the administrator-assigned name for the VLAN.
TYPE	Indicates the type of VLAN, according to the policy used to define its port membership. Options include • byPort—VLAN by port • byIpSubnet—VLAN by IP subnet • byProtocolId—VLAN by protocol ID • bySrcMac—VLAN by source MAC address • byDstMcast—VLAN by destination multicast • bySvlan—VLAN by stacked VLAN • byIds—VLAN by VLAN ID

Field	Description
STG ID	Indicates the STG used by this VLAN to determine the state of its ports. If this VLAN is not associated with an STG, the value is zero.
PROTOCOLID	Indicates the protocol identifier of this VLAN. The options include • none • ip • appleTalk • decLat • decOther • sna802dot2 • snaEthernet2 • netBios • xns • vines • ipV6 • usrDefined • rarp • pPPoE
SUBNETADDR	Indicates the IP subnet address of this VLAN.
SUBNETMASK	Indicates the IP subnet mask of this VLAN.
Vlan Port	
VLAN ID	Indicates the VLAN ID.
PORT MEMBER	Indicates the set of ports that are members (static or dynamic) of this VLAN.
ACTIVE MEMBER	Indicates the set of ports that are currently active in this VLAN. Active ports include all static and dynamic ports.
STATIC MEMBER	Indicates the set of ports that are static members of this VLAN. A static member of a VLAN is always active and never ages out.
NOT_ALLOW MEMBER	Indicates the set of ports that cannot become members of this VLAN.
Vlan ATM Port	
VLAN ID	Indicates the VLAN ID.
PORT NUM	Indicates the port number.
PVC LIST	Indicates the permanent virtual connection (PVC) list.

Field	Description
Ospf Passive Port Members	
VLAN	Indicates the VLAN ID.
PORT NUM	Indicates the VLAN port number for the passive Open Shortest Path First (OSPF) interface.
Vlan Advance	
VLAN ID	Indicates the VLAN ID.
NAME	Indicates the name assigned to the VLAN.
IF INDEX	Indicates the interface index.
QOS LVL	Indicates the Quality of Service (QoS) level packets carried in this VLAN for processing.
AGING TIME	Indicates the timeout period (in seconds) used for aging out dynamic members of this VLAN. This field is only relevant for policy-based VLANs.
MAC ADDRESS	Indicates the media access control (MAC) address assigned to the virtual router interface of this VLAN. This field is meaningful only if rcVlanRoutingEnable is true.
ACTION	Indicates VLAN-related actions. Options include • none—none of the following • flushMacFdb—flush MAC forwarding table • flushArp—flush the ARP table • flushIp—flush the IP route table • flushDynMemb—flush dynamic members • all—flush all tables • flushSnoopMemb—flush the IGMP snoop members • triggerRipUpdate—manually trigger a Routing Information Protocol (RIP) update • flushSnoopMRtr—flush the snoop multicast router
RESULT	Indicates the result from the last VLAN action. Options include • none • inProgress • success • fail
USER DEFINEPED ENCAP	Indicates the encapsulation type for user-defined, protocol-based VLANs.
Vlan Arp	

Field	Description
VLAN ID	Indicates the VLAN ID.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.
NLB-UNIAST-MODE	Indicates the mode for network load balancing (NLB) unicast.
Vlan Fdb	
VLAN ID	Indicates the VLAN ID.
STATUS	Indicates the status of forward database (FDB) forwarding on the VLAN.
MAC ADDRESS	Indicates the MAC address assigned to the virtual router interface of this VLAN. This field is meaningful only if rcVlanRoutingEnable is true.
INTERFACE	Indicates the interface.
MONITOR	Indicates whether monitoring is performed on this unicast MAC address. If monitoring is enabled, a packet received with a matching destination MAC address is forwarded to the port configured to receive monitor traffic.
QOS LVL	Indicates the QoS level packets carried in this VLAN for processing.
SMLT REMOTE	Indicates the MAC address for remote learning, either local or remote.
Vlan Filter	
VLAN ID	Indicates the VLAN ID.
STATUS	Indicates the status of the VLAN filter.
MAC ADDRESS	Indicates the MAC address assigned to the virtual router interface of this VLAN. This field is meaningful only if rcVlanRoutingEnable is true.
PORT	Indicates the port number.
QOS LVL	Indicates the QoS level packets carried in this VLAN for processing.
PCAP	Indicates the status of the Packet Capture (PCAP) tool on the filter.
DEST_DISCARD SET	Indicates a set of ports for traffic arriving from this MAC address.
SRC_DISCARD SET	Indicates a set of ports for traffic arriving and not forwarded to this MAC address.

Field	Description
Vlan Static	
VLAN ID	Indicates the VLAN ID.
STATUS	Indicates the status of the static VLAN.
MAC ADDRESS	Indicates the MAC address assigned to the virtual router interface of this VLAN. This field is meaningful only if rcVlanRoutingEnable is true.
PORT	Indicates the port number.
MONITOR	Indicates whether monitoring is performed on this unicast MAC address. If monitoring is enabled, a packet received with a matching destination MAC address is forwarded to the port configured to receive monitor traffic.
QOS LVL	Indicates the QoS level packets carried in this VLAN for processing.
IDS Vlan Info	
VLAN ID	Indicates the VLAN ID.
MAC LEARNING	Indicates the type of MAC learning.
DISABLED PORTS	Indicates the disabled port numbers.
Vlan Ip	
VLAN ID	Indicates the VLAN ID.
IP ADDRESS	Indicates the IP subnet address of this VLAN. This value is meaningful only if the VLAN type is set to IP subnet. For other VLAN types, the value is 0.0.0.0.
NET MASK	Indicates the IP subnet mask of this VLAN. This value is meaningful only if the VLAN type is set to IP subnet. For other VLAN types, the value is 0.0.0.0.
BCASTADDR FORMAT	Indicates the IP broadcast address format used on this interface.
REASM MAXSIZE	Indicates the size of the largest IP datagram that this entity can reassemble from the incoming IP fragmented datagrams received on this interface.
ADVERTISE WHEN_DOWN	Indicates whether the VLAN state change is notified to Layer 3, provided the VLAN is configured as a routable interface. A VLAN is considered up when • at least one member of the port-based VLAN uses link up • at least one port member of the policy-based VLAN uses an entry in the Multicast Group ID (MGID) • at least one static member of the policy-based VLAN uses link up Otherwise, a VLAN is considered as down. If the value is true, the interface state change is not notifiesd to Layer 3 (that is, it always

Field	Description
	stays up). If the value is false , the VLAN state change is notified to Layer 3 so that the IP-related status reflects the routable interface state.
DIRECTED BROADCAST	Indicates the status of directed broadcast.
RPC	Indicates the status of remote procedure call (RPC).
RPC MODE	Indicates the RPC mode type.
Vlan Dhcp	
VLAN ID	Indicates the VLAN ID number.
IF INDEX	Indicates the interface index number. Numbers 1–256 are ports; numbers greater than 257 are VLANs.
ENABLE	Indicates whether the Dynamic Host Configuration Protocol (DHCP) is enabled on the port.
MAX HOP	Indicates the maximum number of hops a DHCP packet can take from the source device to the destination device (that is, the DHCP client to the DHCP server).
MIN SEC	Indicates the minimum number of seconds to wait between receiving a DHCP packet and forwarding it to the destination device. A value of zero indicates forwarding is performed immediately without delay.
MODE	Indicates what type of DHCP packets this interface supports. A value of none causes all incoming DHCP and BOOTP packets to drop. Options include none, bootp, dhcp, and both.
ALWAYS BCAST	Indicates whether to broadcast DHCP reply packets to the DHCP client on this interface.
Vlan Ospf	
VLAN ID	Indicates the VLAN ID.
ENABLE	Indicates the status of OSPF configured on the port.
HELLO INTERVAL	Indicates the length of time, in seconds (1 to FFFF), between the hello packets that the router sends on the interface.
RTRDEAD INTERVAL	Indicates the number of seconds (1 to FFFF) since the neighbor received hello packets for a router before the neighbor declares the router down.
DESIGRTR PRIORITY	Indicates the priority of this interface; this value is used in multiaccess networks. This field is used in the designated router (DR) election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie in this value occurs, routers use their router ID to elect a DR. The default is 1.

Field	Description
METRIC	Indicates the metric for this type of service (TOS) on this interface. The value of the TOS metric is 10^9 / interface speed. The default is 1. • FFFF—No route exists for this TOS. • POS/IPCP links—Defaults to 0. • 0—The interface speed is the metric value when the state of the interface is up.
AUTHTYPE	Indicates the type of authentication required for the interface. • none—No authentication required. • simple password—All Open Shortest Path First (OSPF) updates received by the interface must contain the authentication key specified in the interface AuthKey field. • Message Digest 5 (MD5) authentication—All OSPF updates received by the interface must contain the MD5 key.
AUTHKEY	Indicates the key (up to eight characters) required when you specify simple password authentication as the interface authentication type.
INTF	Indicates the interface type.
AREA ID	Indicates the area where the host is found. By default, the area that is submitting the OSPF interface is in 0.0.0.0.
Vlan Rip	
PORT NUM	Indicates the ports on the VLAN.
ENABLE	Indicates the status of the Routing Information Protocol (RIP) on the ports for a VLAN.
DEFAULT SUPPLY	Indicates whether to advertise the default route out of this interface. **Important:** The default route is advertised only if it exists in the routing table.
DEFAULT LISTEN	Indicates whether this interface must learn the default route when advertised by another router that connects to the interface.
TRIGGERED UPDATE	Indicates the status of the RIP triggered update on the interface.
AUTOAGG ENABLE	Indicates the status of auto aggregation on the interface.
SUPPLY	Indicates the status of advertising RIP routes through the interface.
LISTEN	Indicates the status of RIP reception on the interface.

Field	Description
POISON	Indicates the status of poison reverse on the interface. If you disable this parameter, split horizon is invoked, meaning that the interface does not advertise IP routes learned from an immediate neighbor back to the neighbor. If you enable this parameter, the RIP update sent to the neighbor from which a route is learned, is poisoned with a metric of 16. The route entry is not passed to the neighbor because, historically, 16 is infinity in terms of hops on a network. The default is disable.
Vlan Vrrp	
VLAN ID	Indicates the VLAN ID.
VRRP ID	Indicates the number which, along with an interface index (ifIndex), serves to uniquely identify a virtual router on a Virtual Router Redundancy Protocol (VRRP) router. A set of one or more associated addresses is assigned to a virtual router ID.
IP ADDR	Indicates the assigned IP addresses that a virtual router backs up.
VIRTUAL MAC ADDR	Indicates the virtual MAC address of the virtual router. This address is derived as follows: 00-00-5E-00-01-<virtual router ID>, where the first three octets consist of the Internet Assigned Numbers Authority (IANA) Organizationally Unique Identifier (OUI), the next two octets indicate the address block of the VRRP protocol, and the remaining octets consist of the virtual router ID.
Vlan Vrrp Extended	
VID	Indicates the VLAN ID.
STATE	Indicates the current state of the virtual router. Options include • initialize—waiting for a startup event • backup—monitoring the state/availability of the master router • master—forwarding IP addresses associated with this virtual router
CONTROL	Indicates the virtual router function. A value of enabled transitions the state of the router from initialize to backup. A value of disabled transitions the router from master or backup to initialize.
PRIORITY	Indicates the priority for the virtual router (for example, master election) with respect to other virtual routers that back up one or more associated IP addresses. Higher values indicate higher priority. A priority of 0 indicates that this router ceased to participate in VRRP, and a backup virtual router transitions to become a new master. A priority of 255 is used for the router that owns the associated IP addresses.

Field	Description
MASTER IPDDR	Indicates the real (primary) IP address of the master router. This address is the IP address listed as the source in the VRRP advertisement last received by this virtual router.
ADVERTISE INTERVAL	Indicates the time interval, in seconds, between sending advertisement messages. Only the master router sends VRRP advertisements.
CRITICAL IPADDR	Indicates the IP address of the interface that causes a shutdown event.
HOLDDOWN_TIME	Indicates the amount of time (in seconds) to wait before preempting the current VRRP master.
ACTION	Indicates the trigger for an action on this VRRP interface. Options include none and preemptHoldDownTimer.
CRITICAL IP ENABLE	Indicates whether a user-defined critical IP address is enabled. No indicates the use of the default IP address (0.0.0.0).
BACKUP MASTER	Indicates the state of designating a backup master router.
BACKUP MASTER STATE	Indicates the state of the backup master router.
FAST ADV INTERVAL	Indicates the faster advertisement interval, in milliseconds, between sending advertisement messages. After you select the faster advertisement interval enable, the faster advertisement interval is used instead of the regular advertisement interval.
FAST ADV ENABLE	Indicates the faster advertisement interval status.
Vlan Ip Igmp	
VLAN ID	Indicates the VLAN ID.
QUERY INTVL	Indicates the interval (in seconds) between IGMP host query packets transmitted on this interface.
QUERY MAX RESP	Indicates the interval (in seconds) for the maximum query response time advertised in IGMPv2 queries on this interface. Smaller values allow a router to prune groups faster.
ROBUST	Indicates the tuning for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that is running on this interface. This object configures a router capable of running either value. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response in a group-specific query.

Field	Description
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop on the VLAN.
SNOOP ENABLE	Indicates the status of IGMP snoop on the VLAN.
SSM SNOOP ENABLE	Indicates the status of SSM IGMP snoop on the VLAN.
FAST LEAVE ENABLE	Indicates the status of fast leave.
FAST LEAVE PORTS	Indicates the ports with fast leave enabled.
Vlan Ip Dvmrp	
IF	Indicates the ifIndex value of the interface for which DVMRP is enabled.
ADDR	Indicates the IP address that this system uses as a source address on this interface.
METRIC	Indicates the distance metric for this interface that is used to calculate distance vectors.
OPERSTAT	Indicates the current operational state of this DVMRP interface.
DEFAULT LISTEN	Indicates whether the switch can learn DVMRP default routes over this interface.
DEFAULT SUPPLY	Indicates if the switch supplies DVMRP default routes over this interface.
DEFAULT METRIC	Indicates the cost of the DVMRP default route that this interface generates and supplies when configured to supply a default route.
ADVERTISE SELF	Indicates whether the switch can advertise this local network.
IN-POLICY	Indicates the DVMRP accept policy name configured on this interface.
OUT-POLICY	Indicates the DVMRP announce policy name configured on this interface.
INTF TYPE	Indicates the type of this DVMRP interface, and whether it uses a tunnel, source routing, a physical interface for which a querier exists, or a physical interface for which a querier does not exist (subnet).
Vlan Ip Icmp Route Discovery	
VLAN ID	Indicates the VLAN ID.
ADV_ADDRESS	Indicates the advertisement address to which the interface transmits route discovery advertisements.
ADV_FLAG	Indicates whether the address is advertised on this interface.

Field	Description
LIFETIME	Indicates the value to place in the lifetime field of a router advertisement sent from the interface.
MAX_INT	Indicates the maximum time allowed between sending router advertisements from this interface.
MIN_INT	Indicates the minimum time allowed between sending router advertisements from this interface.
PREF_LEVEL	Indicates the preferability of the router address as a default router.
Manual Edit Mac	
MAC ADDRESS	Indicates the MAC address that is learned on the port.
PORTS	Indicates the allowed ports that can learn this MAC address.
Autolearn Mac	
MAC ADDRESS	Indicates the MAC address that is automatically learned on the port.
PORT	Indicates the allowed ports that can automatically learn this MAC address.
Vlan Ip Pim	
VLAN-ID	Identifies the VLAN.
PIM-ENABLE	The state of PIM on the VLAN.
MODE	The configured mode of this VLAN. The valid modes are SSM and Sparse.
HELLOINT	Indicates how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Indicates how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default join/prune interval is 60 seconds.
CBSR PREF	The preference for this local interface to become a candidate bootstrap router (C-BSR). The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
INTF TYPE	Indicates whether the PIM interface is active or passive.
Vlan Ip Pgm	
VLAN-ID	Identifies the VLAN.
ENABLE	Displays whether Pragmatic General Multicast (PGM) is enabled or disabled on this interface.
STATE	Indicates the current state (up or down) of PGM.

Field	Description
NAK_RE_XMIT INTERVAL	Specifies how long to wait for a negative acknowledgement confirmation (NCF), in milliseconds, before retransmitting the negative acknowledgement (NAK). The default is 1000 milliseconds.
MAX_NAK_RE XMIT_COUNT	Displays the maximum number of NAK retransmission packets allowed for each second.
NAK_RDATA INTERVAL	Displays how long to wait for retransmitted data (RDATA), in milliseconds, after receiving an NCF.
NAK_ELIMINATE INTERVAL	Displays the length of time (in milliseconds) during which a network element (NE) eliminates duplicate NAKs. After this interval expires, the NE suspends NAK elimination until the first duplicate arrives. After this NAK is forwarded, the NE eliminates duplicate NAKs for the specified interval. You must configure this parameter lower than the NAK_RDATA INTERVAL.
Vlan Mcastmac	
VLAN ID	Indicates the VLAN ID.
MAC ADDRESS	Indicates the MAC address.
PORT LIST	Indicates the list of ports.
MLT GROUPS	Indicates the multilink trunk groups.
Vlan Firewall	
ID	Indicates the VLAN ID.
NAME	Indicates the VLAN name assigned by the user.
FIREWALL TYPE	Indicates the firewall VLAN type for port-based VLANs. Options include • none • naap • enforceable • peering
CLUSTER ID	Indicates the firewall cluster ID.

PGM global information

Use the **`show ip pgm global`** command to display the PGM global status on the switch. The syntax for this command is as follows.

```
show ip pgm global
```

The following table describes the fields for this command.

Table 43: show ip pgm global command

Field	Description
enable	Displays whether PGM is globally enabled or disabled.
state	Displays the current state (up or down) of PGM.
session-life-time	Displays the length of idle time (in seconds) before a session times out. Idle time is when the router does not receive source path messages (SPM) from the upstream router. The default is 300 seconds.
nnak-generate	When enabled, the designated local repairer (DLR) that receives redirected NAKs, where it uses the RDATA, sends a NNAK to the original source.
max-re-xmit-states	Displays the maximum number of retransmit state entries that the switch can create. Each entry uses a unique NAK sequence number. The default is 200 entries.
total-re-xmit-states	Displays the total number of retransmit state entries in the retransmit table.
max-sessions	Displays the maximum number of source path state sessions allowed on the switch. The default is 100 sessions.
total-sessions	Displays the total number of source path state sessions in the PGM session entries table.
total-re-xmit-states-timeout	Displays the total number of retransmit state entries that were removed because they timed out.
total-unique-naks	Displays the total number of unique NAKs received.
total-unique-parity-naks	Displays the total number of unique parity NAKs received.

PIM active RP

Use the **`show ip pim active-rp`** command to display information about the active rendezvous point (RP) for all groups or a specific group. The syntax for this command is as follows.

```
show ip pim active-rp <group> [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Use the data in the following table to use the **show ip pim active-rp** command.

Table 44: Command parameters

Parameter	Value
group	Specifies the IP address of the active RP for a specific group. If you do not specify an IP address, you receive information about the active RP for all the running multicast groups on the switch.

The following table shows the field descriptions for this command.

Table 45: show ip pim active-rp command

Field	Description
GRPADDR	The IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RP-ADDR	The IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.
RP-PRIORITY	The priority of the rendezvous point (RP). C-RPs must send C-RP advertising messages with the field set to 0, which is the highest priority.

PIM bootstrap router

Use the **show ip pim bsr** command to display information about the bootstrap router (BSR) for this PIM-SM domain. The syntax for this command is as follows.

```
show ip pim bsr [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 46: show ip pim bsr command

Field	Description
Current BSR address	Shows the IP address of the current BSR for the local PIM domain.
Current BSR priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
Current BSR HaskMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. With the hash-mask, a small number of consecutive groups (for example, 4) can always hash to the same RP.
Current BSR Fragment	Shows a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same fragment tag.
Pim Boostrap Timer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

PIM interface

Use the **`show ip pim interface`** command to display information about the PIM-SM interface setup on the switch. The syntax of this command is as follows.

```
show ip pim interface [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 47: show ip pim interface command

Field	Description
IF	The slot or port number or VLAN ID of the interface where PIM is enabled.
ADDR	The IP address of the PIM interface.
MASK	The network mask for the IP address of the PIM interface.

Field	Description
MODE	The configured mode of this interface. The valid modes are SSM and Sparse.
DR	Shows the designated router (DR) for this interface.
HLINT	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default join/prune interval is 60 seconds.
CBSPR	The preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OPSTAT	Indicates the status of PIM on this interface: up or down. **Note:** Sometimes the output of the **`show ip pim interface`** command indicates that the operational status (OPSTAT) is down even if the PIM interface is enabled and active. This does not impact traffic. The reason for the discrepancy is that Spanning Tree has not converged. Once the convergence completes, the table updates properly.
INTF TYPE	Indicates whether the PIM interface is active or passive.

PIM mode

Use the **`show ip pim mode`** command to show the PIM mode (SM or SSM). The syntax for this command is as follows.

```
show ip pim mode [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field description for this command.

Table 48: show ip mode command

Field	Description
Mode	Indicates the PIM mode as SM or SSM.

PIM neighbor

Use the **show ip pim neighbor** command to display information about the neighboring routers configured with PIM-SM. The syntax for this command is as follows.

```
show ip pim neighbor [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 49: show ip pim neighbor command

Field	Description
INTERFACE	The interface number.
ADDRESS	The IP address of the PIM neighbor.
UPTIME	The elapsed time since this PIM neighbor last became a neighbor of the local router.
EXPIRE	The time remaining before this PIM neighbor times out.

PIM route

Use the **show ip pim mroute** command to display information from the route table. The syntax for this command is as follows.

```
show ip pim mroute [src <value>] [grp <value>] [vrf <value>] [vrfids
<value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Note:

The **show ip pim mroute** command may display an incorrect incoming port after the corresponding SMLT port is disabled and re-enabled.

The following table shows the field descriptions for this command.

Table 50: show ip pim mnroutecommand

Field	Description
Src	Displays IP address of the source that is sending the multicast stream. A nonzero value indicates that a source is sending multicast traffic. 0.0.0.0 indicates that this entry is created in response to a receiver interested to receive this traffic.
Grp	Displays the IP multicast group address.
RP	Displays the IP address of the RP router.
Upstream	Displays the IP address of the next hop that a multicast packet takes when the switch receives it on the correct port as listed on the incoming interface.
Flags	Displays the flags set based on the condition of the receivers, the RP, and the senders. Use the legend at the bottom of the output to explain the flag values.
Incoming Port	Lists all ports through which a multicast packet can ingress. If a multicast stream enters on another port, it is not processed.
Outgoing Ports	Lists all ports through which traffic that enters on incoming ports egresses.
Joined Ports	Lists all ports that received PIM or IGMP join messages.
Pruned Ports	List all ports that received PIM or IGMP prune messages.
Leaf Ports	Lists multicast receivers directly connected to the router.
Asserted Ports	Lists all ports that received assert messages. The router uses assert messages to help determine the best path to the source.
Prune Pending Ports	Lists all ports currently in the prune-pending state.
Assert Winner Ifs	Lists interfaces elected the assert winner. The winner continues to forward multicast traffic to the LAN.
Assert Loser Ifs	Lists interfaces not elected as the assert winner. The loser interface is pruned.
Timers	Displays the up time and expiration time for the entry in the routing table.

Field	Description
AssertVifTimer	Displays the time after which the assert winner state refreshes.

PIM virtual neighbor

Use the **show ip pim virtual-neighbor** command to display the virtual neighbor. The syntax for this command is as follows.

```
show ip pim virtual-neighbor [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 51: show ip virtual-neighbor command

Field	Description
INTERFACE	Indicates the interface.
ADDRESS	Indicates the IP address of the virtual neighbor.

Rendezvous points

Use the **show ip pim rp-set** command to display information about the RPs for this PIM-SM domain. The syntax for this command is as follows.

```
show ip pim rp-set [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 52: show ip pim rp-set command

Field	Description
GRPADDRESS	The IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.

Field	Description
GRPMASK	The address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
ADDRESS	The IP address of the C-RP router.
COMPONENT	A unique number identifying the protocol instance connected to each PIM domain.
HOLDTIME	The time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
EXPTIME	The time remaining before this C-RP router times out.

SSM channel information

Use the **`show ip igmp ssm-channel`** command to display the list of SSM channels. The syntax for this command is as follows.

```
show ip igmp ssm-channel [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 53: show ip igmp ssm-channel command

Field	Description
GROUP	Indicates the IP multicast group address that uses a default range of 232/8.
SOURCE	Indicates the IP address of the source that sends traffic to the group source.
MODE	Indicates that the entry is a statically configured entry (static) or a dynamically learned entry from IGMPv3 (dynamic).
ACTIVE	Indicates the activity on the corresponding source and group. If the source is active and traffic is flowing to the switch, this field is active; otherwise, it is nonactive.
STATUS	Indicates the administrative state and whether to use the entry. If the value of this field is enabled (default), the entry is used. If the state is disabled, the entry is not used but is saved for future use.

SSM group range and dynamic learning status

Use the **show ip igmp ssm-global** command to display the SSM group range and the status of dynamic learning. The syntax for this command is as follows.

```
show ip igmp ssm-global [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 54: show ip igmp ssm-global command

Field	Description
DYNAMIC LEARNING	Indicates whether dynamic learning is enabled at a global level.
SSM GROUP RANGE	Indicates the IP address range for the SSM group.

Static RP table

Use the **show ip pim static-rp** command to display the static RP table. The syntax for this command is as follows.

```
show ip pim static-rp [vrf <value>] [vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 55: show ip pim static-rp command

Field	Description
GRPADDR	Indicates the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a static RP.
GRPMASK	Indicates the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a static RP.

Field	Description
RPADDR	Indicates the IP address of the static RP. This address must be one of the local PIM-SM enabled interfaces.
STATUS	Indicates the status of static RP.

Static source groups

Use the **show ip mroute static-source-group** command to display information about the static source groups on the current interface. You can see all the valid entries that were created. If an entry is created with an x bit mask, it shows as an x bit in the output. The syntax for this command is as follows.

```
show ip mroute static-source-group [<GroupAddress>] [vrf <value>]
[vrfids <value>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Use the data in the following table to use the **show ip mroute static-source-group** command.

Table 56: Command parameters

Parameter	Value
GroupAddress	Indicates the IP multicast group address.

The following table shows the field descriptions for this command.

Table 57: show ip mroute static-source-group command

Field	Description
Group Address	Indicates the IP multicast group address.
Source Address	Indicates the network address.
Subnet Mask	Indicates the network mask.

VLAN port data

Use the **show vlan info ports** command to display VLAN port data. The syntax for this command is as follows.

```
show vlan info ports [<vid>]
```

Use the data in the following table to use the **show vlan info ports** command.

Table 58: Command parameters

Parameter	Value
vid	Specifies the VLAN ID from 1–4092. The VLAN ID is optional. If you use a VLAN ID, the command displays information for the specified VLAN. Without the VLAN ID, the command displays information for all the configured VLANs.

The following table shows the field descriptions for this command.

Table 59: show vlan info ports command

Field	Description
VLAN ID	Indicates the VLAN ID.
PORT MEMBER	Indicates the set of ports that are members (static or dynamic) of this VLAN.
ACTIVE MEMBER	Indicates the set of ports currently active in this VLAN. Active ports include all static and dynamic ports.
STATIC MEMBER	Indicates the set of ports that are static members of this VLAN. A static member of a VLAN is always active and never ages out.
NOT_ALLOW MEMBER	Indicates the set of ports that cannot become members of this VLAN.
PORT NUM	Indicates the port number.
PVC LIST	Indicates the permanent virtual connection (PVC) list.
VLAN PORT NUM	Indicates the VLAN port number for the passive OSPF interface.

Chapter 36: ACLI show command reference

This reference information provides show commands to view the operational status of multicast routing on the Avaya Ethernet Routing Switch 8800/8600.

ACLI show command reference navigation

- DVMRP on page 548
- DVMRP interface on page 549
- DVMRP neighbor on page 550
- DVMRP route on page 551
- IGMP access on page 552
- IGMP cache on page 553
- IGMP group on page 553
- IGMP interface on page 554
- IGMP multicast router discovery on page 555
- IGMP multicast router discovery neighbors on page 556
- IGMP router-alert on page 557
- IGMP sender on page 557
- IGMP snoop on page 558
- IGMP static and blocked ports on page 558
- Layer 2 multicast MAC filters on page 559
- Layer 3 multicast MAC ARP data on page 559
- Multicast group trace for IGMP snoop on page 560
- Multicast MLT distribution on page 561
- Multicast route information on page 562
- Multicast route next hop on page 563
- Multicast routes on an interface on page 564
- Multicast static route information on page 565

- PGM global information on page 565
- PIM active RP on page 566
- PIM bootstrap router on page 567
- PIM interface on page 568
- PIM mode on page 569
- PIM neighbor on page 569
- PIM route on page 570
- PIM virtual neighbor on page 571
- Rendezvous points on page 572
- SSM channel information on page 573
- SSM group range and dynamic learning status on page 573
- Static RP table on page 574
- Static source groups on page 574
- VLAN port data on page 575

DVMRP

Use the **show ip dvmrp** command to display information about the general Distance Vector Multicast Routing Protocol (DVMRP) group. The syntax for this command is as follows.

```
show ip dvmrp
```

The following table shows the field descriptions for this command.

Table 60: show ip dvmrp command

Field	Description
AdminStat	Indicates the status of DVMRP.
Genid	Indicates the generation identifier for the routing process. This ID is used by neighboring routers to detect whether the DVMRP routing table is present.
Version	Indicates the version of DVMRP.
NumRoutes	Indicates the number of entries in the routing table. This value is used to monitor the routing table size to detect illegal advertisements of unicast routes.
NumReachableRoutes	Indicates the number of entries in the routing table with noninfinite metrics. This value is used to detect

Field	Description
	network partitions by observing the ratio of reachable routes to total routes.
UpdateInterval	Indicates the global route update interval in seconds.
TriggeredUpdateInterval	Indicates the global route triggered update interval in seconds.
LeafTimeOut	Indicates the hold-down timer for leaf in seconds.
NbrTimeOut	Indicates the neighbor timeout interval in seconds.
NbrProbeInterval	Indicates the global neighbor probe interval in seconds.
FwdCacheTimeout	Indicates the timeout interval (in seconds) for aging prune entries in the Forward Cache.
RouteExpireTimeout	Indicates the minimum amount of time remaining before this entry ages out.
RouteDiscardTimeout	Indicates the interval (in seconds) to discard collected routes.
RouteSwitchTimeout	Indicates the interval (in seconds) to discard unused routes.
ShowNextHopTable	Indicates the status of showing the next-hop table.
generate-trap	Indicates the status of DVMRP traps.
generate-log	Indicates the status of logging DVMRP.
PruneResend	Indicates the status of resending prune messages.

DVMRP interface

Use the **`show ip dvmrp interface`** command to display DVMRP route policy information for the DVMRP interface configurations on the switch. You can view information for all interfaces, a specific interface type, or a specific port. The syntax for this command is as follows.

```
show ip dvmrp interface [fastethernet <slot/port>] [gigabitethernet
<slot/port>] [pos <slot/port>] [vlan <vlan_id>]
```

The following table shows the field descriptions for this command.

Table 61: show ip dvmrp interface command

Field	Description
IF	Indicates the ifIndex value of the interface for which DVMRP is enabled.
ADDR	Indicates the IP address this system uses as a source address on this interface.
METRIC	Indicates the distance metric for this interface used to calculate distance vectors.
OPERSTAT	Indicates the current operational state of this DVMRP interface.
DEFAULT LISTEN	Indicates whether the switch can learn DVMRP default routes over this interface.
DEFAULT SUPPLY	Indicates if the switch supplies DVMRP default routes over this interface.
DEFAULT METRIC	Indicates the cost of the DVMRP default route that this interface generates and supplies when configured to supply the default route.
ADVERTISE SELF	Indicates whether the switch can advertise this local network.
IN-POLICY	Indicates the DVMRP accept policy name configured on this interface.
OUT-POLICY	Indicates the DVMRP announce policy name configured on this interface.
INTF TYPE	Indicates the type of this DVMRP interface and whether it uses a tunnel, source routing, a physical interface for which a querier exists, or a physical interface for which a querier does not exist (subnet).

DVMRP neighbor

Use the **`show ip dvmrp neighbor`** command to display information about the configured DVMRP neighbors. The syntax for this command is as follows.

```
show ip dvmrp neighbor
```

The following table shows the field descriptions for this command.

Table 62: show ip dvmrp neighbor command

Field	Description
INTERFACE	Indicates the value of ifIndex for the virtual interface used to reach this DVMRP neighbor.
ADDRESS	Indicates the IP address of the DVMRP neighbor.

Field	Description
EXPIRE	Indicates the minimum time (in seconds) remaining before this DVMRP neighbor ages out.
GENID	Indicates the neighboring router generation identifier.
MAJVER	Indicates the major DVMRP version number of the neighboring router.
MINVER	Indicates the minor DVMRP version number of the neighboring router.
CAPABILITY	Indicates the capabilities of the neighboring router. The leaf bit indicates that the neighbor uses only one interface with its neighbors. The prune bit indicates that the neighbor supports pruning. The generation ID bit indicates that the neighbor sends its generation ID in probe messages. The mtrace bit indicates that the neighbor can handle mtrace requests.
STATE	Indicates the state of DVMRP for the neighboring router.

DVMRP route

Use the **`show ip dvmrp route`** command to display information about the DVMRP routes. The syntax for this command is as follows.

```
show ip dvmrp route
```

The following table shows the field descriptions for this command.

Table 63: show ip dvmrp route command

Field	Description
SOURCE	Indicates the network address that, when combined with the corresponding value of dvmrpRouteSourceMask, identifies the sources for multicast routing information.
MASK	Indicates the network mask that, when combined with the corresponding value of dvmrpRouteSource, identifies the sources for multicast routing information.
UPSTREAM_NBR	Indicates the address of the upstream neighbor from which IP datagrams from these sources are received.
INTERFACE	Indicates the interface on which IP datagrams sent by these sources are received.
METRIC	Indicates the distance in hops to the source subnet.

Field	Description
EXPIRE	Indicates the minimum amount of time remaining (in seconds) before this entry ages out.

IGMP access

Use the **show ip igmp access** command to display information about the Internet Group Management Protocol (IGMP) multicast access control groups. The syntax for this command is as follows.

```
show ip igmp access [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 64: show ip igmp access field descriptions

Field	Description
INTERFACE	Identifies the interface where multicast access control is configured.
GRP PREFIX	An alphanumeric string that identifies the name of the access policy.
HOSTADDR	The IP address of the host.
HOSTMASK	The subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the host network.
ACCESSMODE	Specifies the action of the access policy. The actions are • deny-tx—deny IP multicast transmitted traffic • deny-rx—deny IP multicast received traffic • deny-both—deny both IP multicast transmitted and received traffic • allow-only-rx—allow IP multicast transmitted traffic • allow-only-rx—allow IP multicast received traffic • allow-only-both—allow both IP multicast transmitted and received traffic

IGMP cache

Use the **show ip igmp cache** command to display information about the IGMP cache. The syntax for this command is as follows.

```
show ip igmp cache [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 65: show ip igmp cache command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INTERFACE	Indicates the physical interface or a logical interface (VLAN), which received group reports from various sources.
LASTREPORTER	Indicates the IP address of the source of the last membership report received for this IP multicast group address on this interface. If no membership report is received, this object uses the value 0.0.0.0.
EXPIRATION	Indicates the minimum amount of time remaining before this entry ages out.
V1HOSTTIMER	Indicates the time remaining until the local router assumes that no IGMPv1 members exist on the IP subnet attached to this interface.
TYPE	Indicates whether the entry is learned dynamically or is added statically.
STATICPORTS	Indicates the IP multicast group address for which this entry contains information.

IGMP group

Use the **show ip igmp group** command to display information about the IGMP group. The syntax for this command is as follows.

```
show ip igmp group [count] [group <IP address>] [member-subnet
<default|IP address/mask>] [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 66: show ip igmp group command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INPORT	Indicates the physical interface or a logical interface (VLAN), which received group reports from various sources.
MEMBER	Indicates the IP address of a source that sent a group report to join this group.
EXPIRATION TIME	Indicates the time left before the group report expires on this port. This variable is updated upon receiving a group report.

IGMP interface

Use the **show ip igmp interface** command to display information about the interfaces where IGMP is enabled. This syntax for this command is as follows.

```
show ip igmp interface [fastethernet <slot/port>] [gigabitethernet
<slot/port>] [pos <slot/port>] [vlan <slot/port>] [vrf Word<0-16>]
[vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 67: show ip igmp interface command

Field	Description
IF	Indicates the interface where IGMP is configured.

Field	Description
QUERY INTVL	Indicates the frequency at which the interface transmits IGMP host query packets.
STATUS	Indicates the activation of a row that enables IGMP on the interface. The destruction of a row disables IGMP on the interface.
VERS.	Indicates the version of IGMP that is running on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface is attached.
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received whose IGMP version does not match the igmpInterface version. You must configure all routers on a LAN to run the same version of IGMP. Thus, if the interface receives queries with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times IGMP added a group membership on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.
LASTMEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This does not apply if igmpInterface Version is 1.

IGMP multicast router discovery

Use the **`show ip igmp mrdisc`** command to display information about the IGMP multicast discovery routes. The syntax for this command is as follows.

```
show ip igmp mrdisc [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 68: show ip igmp mrdisc command

Field	Description
VLAN ID	Indicates the VLAN ID.
MRDISC	Indicates the status of multicast router discovery.
DISCOVERED RTR PORTS	Indicates the ports discovered.

IGMP multicast router discovery neighbors

Use the **`show ip igmp mrdisc neighbors`** command to display information about the IGMP multicast router discovery neighbors. The syntax for this command is as follows.

```
show ip igmp mrdisc neighbors [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 69: show ip igmp mrdisc-neighbors command

Field	Description
VLAN ID	Indicates the VLAN ID.
SRC_PORT	Indicates the source port.
IP Addr	Indicates the IP address.
Advert-int	Indicates the advertisement interval in seconds.
QUERY-int	Indicates the query interval in seconds.
Robust-val	Indicates the tuning for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.

IGMP router-alert

Use the **show ip igmp router-alert** command to display the status of IGMP router alert. The syntax for this command is as follows.

```
show ip igmp router-alert [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 70: show ip igmp router-alert command

Field	Description
IFINDEX	Indicates the interface index number.
ROUTER ALERT ENABLE	Indicates the status of the router alert check.

IGMP sender

Use the **show ip igmp sender** command to display information about the IGMP senders. The syntax for this command is as follows.

```
show ip igmp sender [count] [group <IP address>] [member-subnet
<default|IP address/mask>] [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 71: show ip igmp sender command

Field	Description
GRPADDR	Indicates the IP multicast address.
IFINDEX	Indicates the interface index number.
MEMBER	Indicates the IP address of the host.
PORT/MLT	Indicates the IGMP sender ports.

Field	Description
STATE	Indicates if a sender exists because of an IGMP access filter. Options include filtered and nonfiltered.

IGMP snoop

Use the **show ip igmp snooping** command to display the status of IGMP snoop. The syntax of this command is as follows.

```
show ip igmp snooping [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 72: show ip igmp snooping command

Field	Description
IFINDEX	Indicates the interface index number.
SNOOP ENABLE	Indicates the status of IGMP snoop.
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop.
SSM SNOOP ENABLE	Indicates the status of IGMP Source Specific Multicast (SSM) snoop.
STATIC MROUTER PORTS	Indicates the set of ports in this VLAN that provide connectivity to an IP multicast router.
ACTIVE MROUTER PORTS	Indicates the active ports.
MROUTER EXPIRATION TIME	Indicates the multicast querier router aging timeout in seconds.

IGMP static and blocked ports

Use the **show ip igmp static** command to display information about the static and blocked ports for the IGMP-enabled interfaces. The syntax for this command is as follows.

```
show ip igmp static [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 73: show ip igmp static command

Field	Description
GRPADDR	Indicates the IP multicast address. The group address holds the starting range for the address range.
TO-GRPADDR	Indicates the end of the range for the group address.
INTERFACE	Indicates the interface IP address.
STATICPORTS	Indicates the egressing ports.
BLOCKEDPORTS	Indicates the ports not allowed to join.

Layer 2 multicast MAC filters

Use the **`show vlan static-mcastmac`** command to display the Layer 2 multicast media access control (MAC) filters. If you specify a VLAN ID, the command displays information for the specified VLAN. Without the VLAN ID, the command displays information for all the configured VLANs. The syntax for this command is as follows.

```
show vlan static-mcastmac [<vlan 1-4904>]
```

The following table shows the field descriptions for this command.

Table 74: show vlan static-mcastmac command

Field	Description
VLAN ID	Indicates the VLAN ID.
MAC ADDRESS	Indicates the MAC address.
PORT LIST	Indicates the list of ports.
MLT GROUPS	Indicates the MultiLink Trunking (MLT) groups.

Layer 3 multicast MAC ARP data

Use the **`show ip arp static-mcastmac`** command to display Layer 3 multicast MAC ARP data. You can specify optional information to narrow the results to a specific virtual router and

forwarder (VRF) name or ID, or to a specific network and subnet. The syntax for this command is as follows.

```
show ip arp static-mcastmac [-s <IP/subnet value> [vrf <vrf name>]
[vrfids <vrf ID>] [<A.B.C.D]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 75: show ip arp static-mcastmac command

Field	Description
IP_ADDRESS	Indicates the multicast IP address.
MAC ADDRESS	Indicates the multicast MAC address.
VLAN	Indicates the VLAN ID.
PORT LIST	Indicates the list of ports.
MLT ID	Indicates the multilink trunk ID.

Multicast group trace for IGMP snoop

Use the **show ip igmp snoop-trace** command to view multicast group trace information for IGMP snoop. Multicast group trace tracks the data flow path of the multicast streams. This command provides information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port. The syntax for the command is as follows.

```
show ip igmp snoop-trace [source <IP address>] [group <IP address>]
[vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table provides the field descriptions for this command.

Table 76: show ip igmp snoop-trace command

Field	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.

Field	Description
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.
OUT PORT	Indicates the outgoing port number.

Multicast MLT distribution

Use the **`show multicast mlt-distribution`** command to show the multicast MLT distribution configuration. Use the **`show mlt`** command to show multicast flow distribution for each multilink trunk. The syntax for these commands is as follows.

```
show multicast mlt-distribution
```

```
show mlt <1-32> [error {collision|main}] [stats]
```

The following table shows the field descriptions for this command.

Table 77: show mlt command

Field	Description
MLTID	Indicates the multilink trunk ID number.
IFINDEX	Indicates the interface index number.
NAME	Indicates the name of this multilink trunk.
PORT TYPE	Indicates the type of multilink trunk port: access or trunk.
SVLAN TYPE	Indicates the type of multilink trunk port.
MLT ADMIN	Indicates the status of MLT.
MLT CURRENT	Indicates the operational status of MLT.
PORT MEMBERS	Indicates the set of ports that are members of this multilink trunk.
VLAN IDS	Indicates the number of VLANs on the multilink trunk.
MULTICAST DISTRIBUTION	Indicates the status of multicast distribution for each multilink trunk.
NT-STG	Indicates whether this spanning tree group (STG) is operating in Avaya mode or in Cisco mode. • true—Avaya mode • false—Cisco mode

Field	Description
DESIGNATED PORTS	Indicates the designated port for the multilink trunk.
LACP ADMIN	Indicates the administrative status of link aggregation on the multilink trunk.
LACP OPER	Indicates the operational status of link aggregation on the multilink trunk.

Multicast route information

Use the **show ip mroute route** command to display information about the multicast routes set up on the switch. The syntax for this command is as follows.

```
show ip mroute route [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following figure shows sample output for the **show ip mroute route** command. In this table, every stream uses one (*,G) entry and x (S,G) entries, depending on how many servers forward traffic to the same group. The 0.0.0.0 mask is always tied to a (*,G) entry. Every time a new stream comes in, Protocol Independent Multicast (PIM) creates two entries in the table; one is a (*,G) entry that points toward the rendezvous point (RP) router, and the other is an (S,G) entry that points toward the source.

```
================================================================================
                              Mroute Route
================================================================================
GROUP          SOURCE          SRCMASK        UPSTREAM_NBR    IF    EXPIR PROT

224.2.128.119 0.0.0.0         0.0.0.0        104.0.0.207     V504 210   pimsm
224.2.128.119 155.120.50.165 255.255.255.0 104.0.0.207     V504 210   pimsm
224.2.128.119 155.120.51.165 255.255.255.0 155.120.51.204 V155 210   pimsm
224.2.128.119 155.120.51.166 255.255.255.0 155.120.51.204 V155 210   pimsm
224.2.148.113 0.0.0.0         0.0.0.0        104.0.0.208     V504 210   pimsm
224.2.148.113 155.120.50.165 255.255.255.0 104.0.0.208     V504 210   pimsm
224.2.157.11  0.0.0.0         0.0.0.0        104.0.0.208     V504 210   pimsm
224.2.157.11  155.120.50.165 255.255.255.0 104.0.0.207     V504 210   pimsm
224.2.157.11  155.120.51.165 255.255.255.0 155.120.51.204 V155 210   pimsm
224.2.157.11  155.120.51.166 255.255.255.0 155.120.51.204 V155 210   pimsm
```

Figure 26: show ip mroute route output

The following table shows the field descriptions for this command.

Table 78: show ip mroute route command

Field	Description
GROUP	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
SOURCE	Indicates the network address that, when combined with the corresponding value of SRCMASK identifies the sources for which this entry specifies a next hop on an outgoing interface.
SRCMASK	Indicates the network mask that, when combined with the corresponding value of SOURCE identifies the sources for which this entry specifies a next hop on an outgoing interface.
UPSTREAM_NBR	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is unknown.
IF	Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).
EXPIR	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
PROT	Indicates the outgoing mechanism through which this route was learned.

Multicast route next hop

Use the **`show ip mroute next-hop`** command to display information about the next hop for the multicast routes set up on the switch. The syntax for this command is as follows.

```
show ip mroute next-hop [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 79: show ip mroute next-hop command

Field	Description
INTERFACE	Indicates the interface identity.

Field	Description
GROUP	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
SOURCE	Indicates the network address, which when combined with the corresponding value of SRCMASK identifies the sources for which this entry specifies a next hop on an outgoing interface.
SRCMASK	Indicates the network mask, which when combined with the corresponding value of SOURCE identifies the sources for which this entry specifies a next hop on an outgoing interface.
ADDRESS	Indicates the address of the next hop specific to this entry. For most interfaces, this address is identical to GROUP. Non Broadcast Multiple Access (NBMA) interfaces, however, use multiple next-hop addresses on a single outgoing interface.
STATE	Indicates whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. The value forwarding indicates the information is currently used; the value pruned indicates it is not used.
EXPTIME	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
CLOSEHOP	Indicates the minimum number of hops between this router and members of this IP multicast group reached through this next hop on this outgoing interface. IP multicast datagrams for the group that use a TTL less than this number of hops are forwarded to the next hop.
PROTOCOL	Indicates the routing mechanism through which this next hop was learned.

Multicast routes on an interface

Use the **`show ip mroute interface`** command to display information about the multicast routes set up on the switch for a specific interface. The syntax for this command is as follows.

```
show ip mroute interface [<fastethernet|gigabitethernet> <slot/pot>]
[vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 80: show ip mroute interface command

Field	Description
INTERFACE	Indicates the interface.
TTL	Indicates the datagram TTL threshold for the interface. The interface does not forward IP multicast datagrams with a TTL less than this threshold. The default value of 0 means the interface forwards all multicast packets.
PROTOCOL	Indicates the routing protocol running on this interface.

Multicast static route information

Use the **show ip static-mroute** command to view the best route to reach a specific RPF address from either the static multicast or unicast routing table. The syntax for this command is as follows.

```
show ip static-mroute [ip <A.B.C.D>] [rpf <A.B.C.D>] [vrf <WORD 0-
64>] [vrfids <0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

PGM global information

Use the **show ip pgm** command to display the PGM global status on the switch. The syntax for this command is as follows.

```
show ip pgm
```

The following table describes the fields for this command.

Table 81: show ip pgm global command

Field	Description
enable	Displays whether Pragmatic General Multicast (PGM) is globally enabled or disabled.
state	Displays the current state (up or down) of PGM.
session-life-time	Displays the length of idle time (in seconds) before a session times out. Idle time is when the router receives

Field	Description
	no source path messages (SPM) from the upstream router. The default is 300 seconds.
nnak-generate	When enabled, the designated local repairer (DLR) that receives redirected NAKs, where it uses the RDATA, sends a NNAK to the original source.
max-re-xmit-states	Displays the maximum number of retransmit state entries that the switch can create. Each entry uses a unique NAK sequence number. The default is 200 entries.
total-re-xmit-states	Displays the total number of retransmit state entries in the retransmit table.
max-sessions	Displays the maximum number of source path state sessions allowed on the switch. The default is 100 sessions.
total-sessions	Displays the total number of source path state sessions in the PGM session entries table.
total-re-xmit-states-timeout	Displays the total number of retransmit state entries that were removed because they timed out.
total-unique-naks	Displays the total number of unique NAKs received.
total-unique-parity-naks	Displays the total number of unique parity NAKs received.

PIM active RP

Use the **`show ip pim active-rp`** command to display information about the active rendezvous point (RP) for all groups or a specific group. If you do not specify an IP address, you receive information about the active RP for all the running multicast groups on the switch. The syntax for this command is as follows.

```
show ip pim active-rp <group IP address> [vrf Word<0-16>] [vrfids
Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 82: show ip pim active-rp command

Field	Description
GRPADDR	The IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RP-ADDR	The IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.
RP-PRIORITY	The priority of the RP. C-RPs must send C-RP advertising messages with the field set to 0, which is the highest priority.

PIM bootstrap router

Use the **show ip pim bsr** command to display information about the bootstrap router (BSR) for this PIM-SM domain. The syntax for this command is as follows.

```
show ip pim bsr [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 83: show ip pim bsr command

Field	Description
Current BSR address	Shows the IP address of the current BSR for the local PIM domain.
Current BSR priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
Current BSR HaskMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. With the hash-mask, a small number of consecutive groups (for example, 4) can always hash to the same RP.
Current BSR Fragment	Shows a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same fragment tag.

Field	Description
Pim Boostrap Timer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

PIM interface

Use the **show ip pim interface** command to display information about the PIM-SM interface setup on the switch. The syntax of this command is as follows.

```
show ip pim interface [fastEthernet <slot/port>] [gigabitethernet
<slot/port>] [pos <slot/port>] [vlan <slot/port>] [vrf Word<0-16>]
[vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 84: show ip pim interface command

Field	Description
IF	The slot or port number or VLAN ID of the interface where PIM is enabled.
ADDR	The IP address of the PIM interface.
MASK	The network mask for the IP address of the PIM interface.
MODE	The configured mode of this interface. The valid modes are SSM and Sparse.
DR	Shows the designated router (DR) for this interface.
HLINT	Specifies how long to wait (in seconds) before the PIM switch sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM switch sends out the next join/prune message to its upstream neighbors. The default join/prune interval is 60 seconds.
CBSPR	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OPSTAT	Indicates the status of PIM on this interface: up or down.

Field	Description
	Note: Sometimes the output of the **show ip pim interface** command indicates that the operational status (OPSTAT) is down even if the PIM interface is enabled and active. This does not impact traffic. The reason for the discrepancy is that Spanning Tree has not converged. Once the convergence completes, the table updates properly.
INTF TYPE	Indicates whether the PIM interface is active or passive.

PIM mode

Use the **show ip pim mode** command to show the PIM mode (SM or SSM). The syntax for this command is as follows.

```
show ip pim mode [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field description for this command.

Table 85: show ip mode command

Field	Description
Mode	Indicates the PIM mode as SM or SSM.

PIM neighbor

Use the **show ip pim neighbor** command to display information about the neighboring routers configured with PIM-SM. The syntax for this command is as follows.

```
show ip pim neighbor [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 86: show ip pim neighbor command

Field	Description
INTERFACE	The interface number.
ADDRESS	The IP address of the PIM neighbor.
UPTIME	The elapsed time since this PIM neighbor last became a neighbor of the local router.
EXPIRE	The time remaining before this PIM neighbor times out.

PIM route

Use the **show ip pim mroute** command to display information from the route table. The syntax for this command is as follows.

```
show ip pim mroute [group <IP address>] [source <IP address>] [vrf
Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

Note:

The **show ip pim mroute** command may display an incorrect incoming port after the corresponding SMLT port is disabled and re-enabled.

The following table shows the field descriptions for this command.

Table 87: show ip pim mnroute command

Field	Description
Src	Displays IP address of the source that is sending the multicast stream. A nonzero value indicates that a source is sending multicast traffic. 0.0.0.0 indicates that this entry is created in response to a receiver interested to receive this traffic.
Grp	Displays the IP multicast group address.
RP	Displays the IP address of the RP router.
Upstream	Displays the IP address of the next hop that a multicast packet takes when the switch

Field	Description
	receives it on the correct port as listed on the incoming interface.
Flags	Displays the flags set based on the condition of the receivers, the RP, and the senders. Use the legend at the bottom of the output to explain the flag values.
Incoming Port	Lists all ports through which a multicast packet can ingress. If a multicast stream enters on another port, it is not processed.
Outgoing Ports	Lists all ports through which traffic entering on incoming ports are egressed.
Joined Ports	Lists all ports that received PIM or IGMP join messages.
Pruned Ports	List all ports that received PIM or IGMP prune messages.
Leaf Ports	Lists multicast receivers directly connected to the router.
Asserted Ports	Lists all ports that received assert messages. The router uses assert messages to help determine the best path to the source.
Prune Pending Ports	Lists all ports currently in the prune-pending state.
Assert Winner Ifs	Lists interfaces elected the assert winner. The winner continues to forward multicast traffic to the LAN.
Assert Loser Ifs	Lists interfaces not elected as the assert winner. The loser interface is pruned.
Timers	Displays the up time and expiration time for the entry in the routing table.
AssertVifTimer	Displays the time after which the assert winner state refreshes.

PIM virtual neighbor

Use the **`show ip pim virtual-neighbor`** command to display the virtual neighbor. The syntax for this command is as follows.

```
show ip pim virtual-neighbor [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 88: show ip virtual-neighbor command

Field	Description
INTERFACE	Indicates the interface.
ADDRESS	Indicates the IP address of the virtual neighbor.

Rendezvous points

Use the **show ip pim rp-hash** command to display information about the RPs for this PIM-SM domain. The syntax for this command is as follows.

```
show ip pim rp-hash [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 89: show ip pim rp-hash command

Field	Description
GRPADDRESS	The IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	The address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
ADDRESS	The IP address of the C-RP router.
HOLDTIME	The time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
EXPTIME	The time remaining before this C-RP router times out.

SSM channel information

Use the **show ip igmp ssm-map** command to display the list of SSM channels. The syntax for this command is as follows.

```
show ip igmp ssm-map [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 90: show ip igmp ssm-map command

Field	Description
GROUP	Indicates the IP multicast group address that uses default range of 232/8.
SOURCE	Indicates the IP address of the source that sends traffic to the group source.
MODE	Indicates that the entry is a statically configured entry (static) or a dynamically learned entry from IGMPv3 (dynamic).
ACTIVE	Indicates the activity on the corresponding source and group. If the source is active and traffic is flowing to the switch, this field status is active; otherwise, it is nonactive.
STATUS	Indicates the administrative state and whether to use the entry. If the value of this field is enabled (default), the entry is used. If the state is disabled, the entry is not used but is saved for future use.

SSM group range and dynamic learning status

Use the **show ip igmp ssm** command to display the SSM group range and the status of dynamic learning. The syntax for this command is as follows.

```
show ip igmp ssm [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 91: show ip igmp ssm command

Field	Description
DYNAMIC LEARNING	Indicates whether dynamic learning is enabled at a global level.
SSM GROUP RANGE	Indicates the IP address range for the SSM group.

Static RP table

Use the **show ip pim static-rp** command to display the static RP table. The syntax for this command is as follows.

```
show ip pim static-rp [vrf Word<0-16>] [vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 92: show ip pim static-rp command

Field	Description
GRPADDR	Indicates the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a static RP.
GRPMASK	Indicates the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a static RP.
RPADDR	Indicates the IP address of the static RP. This address must be one of the local PIM-SM enabled interfaces.
STATUS	Indicates the status of static RP.

Static source groups

Use the **show ip mroute static-source-group** command to display information about the static source groups on the current interface. You can see all the valid entries that were created. If an entry is created with a x bit mask, it shows as a x bit in the output. The syntax for this command is as follows.

```
show ip mroute static-source-group [<GroupAddress>] [vrf Word<0-16>]
[vrfids Word<0-255>]
```

If you do not specify a VRF name or range of VRF IDs, the results show information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results show information only for the VRFs you specify.

The following table shows the field descriptions for this command.

Table 93: show ip mroute static-source-group command

Field	Description
Group Address	Indicates the IP multicast group address.
Source Address	Indicates the network address.
Subnet Mask	Indicates the network mask.

VLAN port data

Use the **show vlan members** command to display VLAN port data. The syntax for this command is as follows.

```
show vlan members [<vlan id>]
```

The following table shows the field descriptions for this command.

Table 94: show vlan members command

Field	Description
VLAN ID	Indicates the VLAN ID.
PORT MEMBER	Indicates the set of ports that are members (static or dynamic) of this VLAN.
ACTIVE MEMBER	Indicates the set of ports currently active in this VLAN. Active ports include all static and dynamic ports.
STATIC MEMBER	Indicates the set of ports that are static members of this VLAN. A static member of a VLAN is always active and is never aged out.
NOT_ALLOW MEMBER	Indicates the set of ports that cannot become members of this VLAN.
PORT NUM	Indicates the port number.
PVC LIST	Indicates the PVC list.
VLAN PORT NUM	Indicates the VLAN port number for the passive OSPF interface.

Chapter 37: Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Navigation

- Getting technical documentation on page 577
- Getting Product training on page 577
- Getting help from a distributor or reseller on page 577
- Getting technical support from the Avaya Web site on page 578

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting Product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.