AVAYA

Configuring IP Routing and Multicast on Avaya Ethernet Routing Switch 4000 Series

Release 5.7
NN47205-506
Issue 05.01
November 2013

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 generally makes available to users of its products. 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 its hardware and Software ("Product(s)"). 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. 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. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products or pre-installed on hardware products, and any upgrades, updates, bug fixes, or modified versions.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED 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 AUTHORIZED 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").

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users.

Licence types

Designated System(s) License (DS). End User may install and use each copy 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, 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". 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, or hardware provided by Avaya. All content on this site, the documentation 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 that 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"). 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. 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. 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 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 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.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product notices and articles, or to report a problem with your Avaya product. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com, scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Chapter 1: Introduction 15
- Purpose 15
- Related resources 15
- Support 17

Chapter 2: New in this release 19
- Features 19
- Other changes 19

Chapter 3: IP routing fundamentals 21
- ACLI command modes 21
- IP addressing overview 22
  - Subnet addressing 24
- IP routing 25
  - IP routing using VLANs 26
  - Local routes 26
  - Static routes 27
  - Non-local static routes 28
  - Static routes 28
  - Default routes 29
  - Route scaling 30
  - Dynamic Routing Table Allocation 30
  - Management VLAN 31
- Brouter port 32
- Related routing features 33
  - DHCP relay 33
  - UDP broadcast forwarding 36
  - Directed broadcasts 37
  - ARP 38
  - Static ARP 38
  - Proxy ARP 39
  - IP blocking for stacks 39
- Open Shortest Path First (OSPF) protocol 41
  - Overview 41
  - Autonomous system and areas 42
  - OSPF neighbors 43
  - Designated routers 44
  - OSPF Operation 44
  - OSPF route advertisements 45
  - Router types 45
  - LSA types 46
  - Area types 47
  - Area aggregation 49
  - SPF calculation 50
  - OSPF virtual link 51
  - OSPF host route 52

OSPF interfaces.... 52
OSPF packets.... 53
OSPF metrics.... 54
OSPF security mechanisms.... 54
Routing Information Protocol.... 55
RIP Operation.... 55
RIP metrics.... 56
RIP routing updates.... 57
RIP configuration.... 57
RIP Features.... 58
Virtual Router Redundancy Protocol.... 58
Equal Cost Multi Path (ECMP).... 61
Route Policies.... 62
Route policies in a stack.... 63
Chapter 4: IGMP fundamentals.... 65
Overview of IP multicast.... 65
Multicast groups.... 66
Multicast addresses.... 67
IGMP overview.... 67
IGMPv1 operation.... 68
IGMPv2 operation.... 68
IGMP requests for comment.... 70
IGMP snooping.... 70
IGMP proxy.... 72
Forwarding of reports.... 73
Static mrouter port and nonquerier.... 74
Unknown multicast filtering.... 74
Robustness value.... 75
IGMP snooping configuration rules.... 75
Default IGMP values.... 76
IGMP snooping interworking with Windows clients.... 77
IGMPv3 snooping.... 77
IGMP Querier.... 78
IGMP Selective Channel Block.... 78
Chapter 5: IP routing configuration using ACLI.... 81
Configuring global IP routing status.... 81
Displaying global IP routing status.... 82
Dynamic Routing Table allocation configuration using ACLI.... 82
Configuring Dynamic Routing Total Routes using ACLI.... 82
Configuring Dynamic Routing Total Routes to default using ACLI.... 83
Viewing Dynamic Routing Total Routes information using ACLI.... 84
Configuring an IP address for a VLAN.... 84
Configuring IP routing status on a VLAN.... 85
Displaying the IP address configuration and routing status for a VLAN.... 86
Displaying IP routes.... 87
Brouter port configuration using ACLI.... 88
Configuring a brouter port using ACLI.... 88

Displaying the brouter port configuration using ACLI ... 89
Modifying the brouter port IP address using ACLI ... 90
Deleting the brouter port using ACLI ... 90
Disabling IP routing for the brouter port using ACLI ... 91
Chapter 6: Static route configuration using ACLI ... 93
Configuring a static route ... 93
Displaying static routes ... 94
Configuring a management route ... 95
Displaying the management routes ... 96
Chapter 7: OSPF configuration using ACLI ... 97
Prerequisites ... 97
Enabling OSPF globally ... 97
Configuring the router ID ... 98
Configuring the OSPF default cost metric ... 99
Configuring OSPF RFC 1583 compatibility ... 99
Configuring the OSPF hold down timer ... 100
Enabling OSPF system traps ... 101
Displaying global OSPF parameters ... 102
Configuring OSPF area parameters ... 102
Displaying OSPF area configuration ... 103
Displaying OSPF area range information ... 104
Enabling OSPF on an IP interface ... 105
Assigning an interface to an OSPF area ... 106
Configuring OSPF for an interface ... 106
Displaying OSPF interface timers ... 108
Displaying OSPF timers for virtual links ... 109
Displaying OSPF interface configurations ... 110
Displaying OSPF neighbors ... 110
Specifying a router as an ASBR ... 111
Configuring the OSPF authentication type for an interface ... 112
Configuring simple authentication keys for OSPF interfaces ... 112
Defining MD5 keys for OSPF interfaces ... 113
Displaying OSPF MD5 keys ... 114
Applying an MD5 key to an OSPF interface ... 115
Displaying OSPF interface authentication configuration ... 116
Configuring a virtual link ... 116
Creating a virtual interface message digest key ... 118
Enabling automatic virtual links ... 118
Displaying OSPF virtual links ... 121
Displaying OSPF virtual neighbors ... 122
Configuring an OSPF host route ... 122
Displaying OSPF host routes ... 123
Displaying the OSPF link state database ... 124
Displaying the external link state database ... 125
Initiating an SPF run to update the OSPF LSDB ... 125
Displaying OSPF default port metrics ... 126
Displaying OSPF statistics ... 126

Displaying OSPF interface statistics.... 127
Clearing OSPF statistics counters.... 128
Chapter 8: OSPF configuration examples using ACLI.... 129
Basic OSPF configuration examples.... 129
Basic OSPF configuration.... 129
Basic ASBR configuration.... 130
Configuring ECMP for OSPF using ACLI.... 131
Setting the number of ECMP paths using ACLI.... 131
Displaying the ECMP configuration using ACLI.... 132
Advanced OSPF configuration examples.... 133
Configuring an IP OSPF interface.... 133
OSPF security configuration example using Message Digest 5.... 134
Configuring OSPF network types.... 135
Configuring Area Border Routers (ABR).... 135
Configuring Autonomous System Border Routers (ASBR).... 138
Stub area configuration example.... 140
NSSA configuration example.... 142
Controlling NSSA external route advertisements.... 144
Configuring a multi-area complex.... 146
Diagnosing neighbor state problems.... 176
Chapter 9: RIP configuration using ACLI.... 179
Prerequisites.... 179
Enabling RIP globally.... 179
Configuring global RIP timers.... 180
Configuring the default RIP metric value.... 181
Displaying global RIP information.... 182
Configuring RIP on an interface.... 183
Displaying the global RIP configuration.... 185
Displaying RIP interface configuration.... 186
Manually triggering a RIP update.... 187
Chapter 10: RIP configuration examples using ACLI.... 189
RIP configuration tasks.... 189
Configuring RIP.... 190
Configuring RIP version 2.... 194
Using RIP accept policies.... 195
Using RIP announce policies.... 197
Chapter 11: VRRP configuration using ACLI.... 199
Configuring global VRRP status using ACLI.... 199
Assigning an IP address to a virtual router ID using ACLI.... 200
Assigning the router priority for a virtual router ID using ACLI.... 200
Configuring the status of the virtual router using ACLI.... 201
Configuring the VRRP critical IP address using ACLI.... 201
Configuring the VRRP critical IP status using ACLI.... 202
Configuring the VRRP holddown timer using ACLI.... 202
Configuring the VRRP holddown action using ACLI.... 203
Configuring the VRRP advertisement interval using ACLI.... 204
Configuring the VRRP fast advertisement interval using ACLI.... 204

Configuring the VRRP fast advertisement status using ACLI.... 205
Configuring ICMP echo replies using ACLI.... 205
Displaying VRRP configuration information using ACLI.... 206
VRRP configuration example 1.... 207
VRRP configuration example 2.... 211
Chapter 12: ECMP configuration using ACLI.... 215
Configuring the number of ECMP paths allotted for RIP using ACLI.... 215
Configuring the number of ECMP paths for OSPF using ACLI.... 216
Configuring the number of ECMP paths for static routes using ACLI.... 217
Displaying global ECMP path information using ACLI.... 218
ECMP configuration examples.... 218
Chapter 13: Route policies configuration using ACLI.... 221
Route policies configuration navigation.... 221
Configuring prefix lists.... 221
Configuring route maps.... 222
Displaying route maps.... 225
Applying a RIP accept (in) policy.... 225
Applying a RIP announce (out) policy.... 226
Configuring an OSPF accept policy.... 227
Applying the OSPF accept policy.... 228
Displaying the OSPF accept policy.... 229
Configuring an OSPF redistribution policy.... 229
Applying the OSPF redistribution policy.... 231
Displaying the OSPF redistribution policy.... 231
Chapter 14: DHCP relay configuration using ACLI.... 233
Prerequisites to DHCP relay configuration using ACLI.... 233
Configuring global DHCP relay status using ACLI.... 233
Displaying the global DHCP relay status using ACLI.... 234
Specifying a local DHCP relay agent and remote DHCP server using ACLI.... 234
Displaying the DHCP relay global configuration using ACLI.... 236
Configuring the maximum packet length for DHCP relay using the ACLI.... 236
Configuring Option 82 for DHCP relay globally using the ACLI.... 237
Assigning an Option 82 for DHCP Relay subscriber Id to a port using the ACLI.... 238
Configuring DHCP relay on a VLAN using ACLI.... 239
Displaying the DHCP relay configuration for a VLAN using ACLI.... 240
Displaying the DHCP relay configuration for a port using ACLI.... 241
Displaying DHCP relay counters using ACLI.... 242
Clearing DHCP relay counters for a VLAN using ACLI.... 243
Chapter 15: UDP broadcast forwarding configuration using ACLI.... 245
Prerequisites to UDP broadcast forwarding using ACLI.... 245
UDP broadcast forwarding configuration procedures.... 245
Configuring UDP protocol table entries.... 246
Displaying the UDP protocol table.... 246
Configuring a UDP forwarding list.... 247
Applying a UDP forwarding list to a VLAN.... 248
Displaying the UDP broadcast forwarding configuration.... 249
Clearing UDP broadcast counters on an interface.... 250

Chapter 16: Directed broadcasts configuration using ACLI 253
Configuring directed broadcasts 253
Displaying the directed broadcast configuration 253
Chapter 17: Static ARP and Proxy ARP configuration using ACLI 255
Static ARP configuration 255
Configuring a static ARP entry 255
Displaying the ARP table 256
Displaying ARP entries 256
Configuring a global timeout for ARP entries 257
Restoring default timeout for ARP entries using ACLI 258
Clearing the ARP cache 259
Proxy ARP configuration 259
Configuring proxy ARP status 259
Displaying proxy ARP status on a VLAN 260
Chapter 18: IP blocking configuration using ACLI 261
Configuring IP blocking for a stack 261
Displaying IP blocking status 261
Chapter 19: IGMP snooping configuration using ACLI 263
Displaying the switch IGMP snooping configuration status using ACLI 263
Displaying IGMP interface information using ACLI 264
Creating an IGMP VLAN interface using ACLI 266
Deleting an IGMP VLAN interface using ACLI 266
Enabling or disabling IGMP snooping for a VLAN using ACLI 267
Adding static mrouter ports to a VLAN using ACLI 267
Removing static mrouter ports from a VLAN using ACLI 268
Enabling or disabling IGMP proxy on a VLAN using ACLI 269
Configuring IGMP snooping robustness for a VLAN using ACLI 270
Configuring the IGMP last member query interval for a VLAN using ACLI 270
Configuring the IGMP query interval for a VLAN using ACLI 271
Configuring the IGMP maximum query response time for a VLAN using ACLI 272
Enabling or disabling IGMP send query on a VLAN using ACLI 273
Configuring the IGMP version on a VLAN using ACLI 273
Enabling or disabling IGMP router alert on a VLAN using ACLI 274
Displaying IGMP router alert configuration information using ACLI 275
Applying the IGMP filter profile on an Ethernet interface using ACLI 276
Clearing IGMP profile statistics using ACLI 276
Enabling or disabling unknown multicast flooding using ACLI 277
Displaying IGMP profiles using ACLI 278
Configuring an IGMP profile using ACLI 279
Deleting an IGMP profile using ACLI 280
Enabling or disabling an IGMP profile on a port 280
Displaying IGMP cache information using ACLI 281
Displaying IGMP group information using ACLI 282
Displaying extended IGMP group information using ACLI 283
Displaying VLAN multicast address flooding information using ACLI 285
Configuring multicast VLAN flooding using ACLI 285
Displaying the VLAN unknown multicast no-flood status using ACLI 287

Flushing the IGMP router table using ACLI 287
Deleting an IGMP filter profile from an Ethernet interface using ACLI 288
SSM configuration using ACLI 289
Displaying the SSM for IGMP configuration using ACLI 289
Configuring SSM for IGMP using ACLI 289
Disabling SSM for IGMP using ACLI 290
Setting SSM for IGMP to default using ACLI 290
SSM map configuration using ACLI 291
Displaying the SSM mapping table using ACLI 291
Configuring an SSM map for IGMP using ACLI 292
Disabling IGMP SSM mapping using ACLI 292
Job aid: Roadmap of IGMP ACLI commands 293
Chapter 20: IP routing configuration using Enterprise Device Manager 297
Configuring routing globally using EDM 297
Viewing VLAN IP Addresses using EDM 299
Displaying IP routes using EDM 299
Configuring ECMP using EDM 301
Configuring a brouter port using EDM 302
Chapter 21: Static route configuration using Enterprise Device Manager 303
Prerequisites 303
Configuring static routes using EDM 303
IP route information display using EDM 304
Viewing IP routes using EDM 305
Filtering IP route information using EDM 306
Viewing TCP information for the switch using EDM 307
Viewing TCP connections using EDM 308
Viewing TCP Listeners using EDM 308
Viewing UDP endpoints using EDM 309
Chapter 22: OSPF configuration using Enterprise Device Manager 313
Configuring OSPF globally using EDM 313
Configuring an OSPF area using EDM 315
Configuring an area aggregate range using EDM 317
Configuring OSPF stub area metrics using EDM 318
Configuring OSPF interfaces using EDM 319
Configuring OSPF interface metrics using EDM 321
Defining MD5 keys for OSPF interfaces 322
Displaying OSPF neighbor information 322
Configuring an OSPF virtual link using EDM 324
Defining MD5 keys for OSPF virtual links using EDM 325
Displaying virtual neighbor information using EDM 326
Configuring OSPF host routes using EDM 327
Displaying link state database information using EDM 328
Displaying external link state database information using EDM 329
Displaying OSPF statistics using EDM 330
Chapter 23: RIP configuration using Enterprise Device Manager 333
Configuring global RIP properties using EDM 333
Configuring a RIP interface using EDM 334

Configuring advanced RIP interface properties using EDM.......... 336
Displaying RIP statistics using EDM.......... 337
Chapter 24: VRRP configuration using Enterprise Device Manager.......... 339
Assigning a virtual router IP address using EDM.......... 339
Configuring VRRP globally using EDM.......... 340
Configuring VRRP interfaces using EDM.......... 341
Graphing VRRP interface information using EDM.......... 343
Viewing general VRRP statistics using EDM.......... 344
Chapter 25: DHCP relay configuration using Enterprise Device Manager.......... 347
Configuring global DHCP Relay using EDM.......... 347
Configuring DHCP Relay using EDM.......... 348
Configuring DHCP Relay with Option 82 for a VLAN using EDM.......... 349
Variable definitions.......... 349
Assigning an Option 82 for DHCP Relay subscriber Id to a port using EDM.......... 350
Viewing and graphing DHCP counters on a VLAN using EDM.......... 351
Chapter 26: UDP broadcast forwarding configuration using Enterprise Device Manager.......... 353
Configuring UDP protocol table entries using EDM.......... 353
Configuring UDP forwarding entries using EDM.......... 354
Configuring a UDP forwarding list using EDM.......... 355
Applying a UDP forwarding list to a VLAN using EDM.......... 356
Chapter 27: Static ARP and Proxy ARP configuration using Enterprise Device Manager.......... 359
Configuring static ARP entries using EDM.......... 359
Configuring proxy ARP using EDM.......... 360
Chapter 28: ECMP configuration using Enterprise Device Manager.......... 363
Variable definitions.......... 363
Chapter 29: Route policies configuration using Enterprise Device Manager.......... 365
Creating a prefix list using EDM.......... 365
Creating a route policy using EDM.......... 366
Configuring RIP in and out policies using EDM.......... 369
Configuring an OSPF Accept Policy using EDM.......... 370
Configuring OSPF redistribution parameters using EDM.......... 371
Applying an OSPF accept or redistribution policy using EDM.......... 372
Chapter 30: IGMP snooping configuration using Enterprise Device Manager.......... 375
Displaying VLAN IGMP group information using EDM.......... 375
Enabling or disabling unknown multicast flooding using EDM.......... 376
Multicast MAC address flooding using EDM.......... 376
Multicast IP address flooding using EDM.......... 378
Configuring SSM for IGMP using EDM.......... 381
SSM map configuration using EDM.......... 382
IGMP interface configuration using EDM.......... 384
IGMP snooping configuration for interfaces using EDM.......... 391
Displaying interface IGMP group information using EDM.......... 394
Displaying extended interface IGMP group information using EDM.......... 394
Displaying IGMP cache information using EDM.......... 395
IGMP profile configuration using EDM.......... 396

Configuring an IGMP profile range using EDM........ 398
Chapter 31: IP Routing capabilities and limitations........ 401
Glossary........ 403

Chapter 1: Introduction

Purpose

This document provides procedures and conceptual information to configure IP routing features on the Avaya Ethernet Routing Switch 4000 Series, including static routes, Proxy ARP, DHCP Relay, and UDP forwarding. It also provides procedures and conceptual information to manage multicast traffic using IGMP snooping.

Related resources

Documentation

For a list of the documentation for this product, see *Documentation Reference for Avaya Ethernet Routing Switch 4000 Series*, NN47205–101.

Training

Ongoing product training is available. For more information or to register, see http://avaya-learning.com/.

Enter the course code in the **Search** field and click **Go** to search for the course.

Course code	Course title
8D00020E	Stackable ERS and VSP Products Virtual Campus Offering

Avaya Mentor videos

Avaya Mentor is an Avaya-run channel on YouTube that includes technical content on how to install, configure, and troubleshoot Avaya products.

Go to http://www.youtube.com/AvayaMentor and perform one of the following actions:

- Enter a key word or key words in the Search Channel to search for a specific product or topic.
- Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the site.

Searching a document collection

On the Avaya Support website, you can download the documentation library for a specific product and software release to perform searches across an entire document collection. For example, you can perform a single, simultaneous search across the collection to quickly find all occurrences of a particular feature. Use this procedure to perform an index search of your documentation collection.

Before you begin

- Download the documentation collection zip file to your local computer.
- You must have Adobe Acrobat or Adobe Reader installed on your computer.

Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named <*product_name_release*>.pdx, for example, ers4000_5.7x.pdx.
3. In the Search dialog box, select the option **In the index named <*product_name_release*>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
   - Whole words only
   - Case-Sensitive
   - Include Bookmarks
   - Include Comments
6. Click **Search**.
   The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance ranking.

Support

Visit the Avaya Support website at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Chapter 2: New in this release

The following sections detail what is new in *Configuring IP Routing and Multicast on Avaya Ethernet Routing Switch 4000 Series*, NN47205-506 for Release 5.7.

Features

There are no new features in this release.

Other changes

See the following section for information about changes that are not feature-related.

New Introduction chapter

The Introduction chapter replaces the Purpose of this document and Customer service chapters.

IGMP profiling on an interface port

Use this feature to apply an IGMP profile on an interface port.

For more information, see Enabling or disabling an IGMP profile on a port on page 280.

Chapter 3: IP routing fundamentals

This chapter provides an introduction to IP routing and related features used in the Avaya Ethernet Routing Switch 4000 Series.

ACLI command modes

ACLI provides the following command modes:

- User EXEC
- Privileged EXEC
- Global Configuration
- Interface Configuration
- Router Configuration
- Application Configuration

Mode access is determined by access permission levels and password protection.

If no password is set, you can enter ACLI in User EXEC mode and use the **`enable`** command to move to the next level (Privileged EXEC mode). However, if you have read-only access, you cannot progress beyond User EXEC mode, the default mode. If you have read-write access you can progress from the default mode through all of the available modes.

With sufficient permission, you can use the rules in the following table to move between the command modes.

Table 1: ACLI command modes

Command mode and sample prompt	Entrance commands	Exit commands

User EXEC
`4548GT-PWR>`
	No entrance command, default mode	`exit`
		or
		`logout`

Privileged EXEC
`4548GT-PWR#`
	`enable`	`exit`
		or
		`logout`

Global Configuration `4548GT-PWR(config)#`	`configure terminal`	mode, enter: `end`

Command mode and sample prompt	Entrance commands	Exit commands

		or
		`exit`
		To exit ACLI completely, enter:
		`logout`

Interface Configuration
`4548GT-PWR(config-if)#`
You can configure the following interfaces:
• Ethernet
• VLAN
	From Global Configuration mode: To configure a port, enter: `interface ethernet <port number>`
	To configure a VLAN, enter: `interface vlan <vlan number>`
		To return to Global Configuration mode, enter:
		`Exit`
		To return to Privileged EXEC mode, enter:
		`end`
		To exit ACLI completely, enter:
		`logout`

Router Configuration `4548GT-(configrouter)#` You can configure the following routers: • RIP • OSPF • VRRP • ISIS	From Global or Interface Configuration mode: To configure RIP, enter `router rip`. To configure OSPF, enter `router ospf`. To configure VRRP, enter `router vrrp`. To configure IS-IS, enter `router isis`.	To return to Global Configuration mode, enter `exit`. To return to Privileged EXEC mode, enter `end`. To exit ACLI completely, enter `logout`.
Application Configuration `4850GT-(config-app)`	From Global, Interface or Router Configuration mode, enter `application`.	To return to Global Configuration mode, enter `exit`. To return to Privileged EXEC mode, enter `end`. To exit ACLI completely, enter `logout`.

IP addressing overview

An IP version 4 (IPv4) address consists of 32 bits expressed in a dotted-decimal format (XXX.XXX.XXX.XXX). The IPv4 address space is divided into classes, with classes A, B, and C reserved for unicast addresses, and accounting for 87.5 percent of the 32-bit IP address space. Class D is reserved for multicast addressing. The following table lists the breakdown of the IP address space by address range and mask.

Table 2: IP address classifications

Class	Address Range	Mask	Number of Networks	Nodes per Network
A	1.0.0.0 - 127.0.0.0	255.0.0.0	127	16 777 214
B	128.0.0.0 - 191.255.0.0	255.255.0.0	16 384	65 534
C	192.0.0.0 - 223.255.255.0	255.255.255.0	2 097 152	255
D	224.0.0.0 - 239.255.255.254			
E	240.0.0.0 - 240.255.255.255			

Note:

Class D addresses are primarily reserved for multicast operations, although the addresses 224.0.0.5 and 224.0.0.6 are used by OSPF and 224.0.0.9 is used by RIP.

Note:

Although technically part of Class A addressing, network 127 is reserved for loopback.

Note:

Class E addresses are reserved for research purposes.

To express an IP address in dotted-decimal notation, each octet of the IP address is converted to a decimal number and separated by decimal points. For example, the 32-bit IP address 10000000 00100000 00001010 10100111 is expressed in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary notation, has a different boundary point between the network and host portions of the address, as shown in the following figure. The network portion is a network number field from 8 through 24 bits. The remaining 8 through 24 bits identify a specific host on the network.

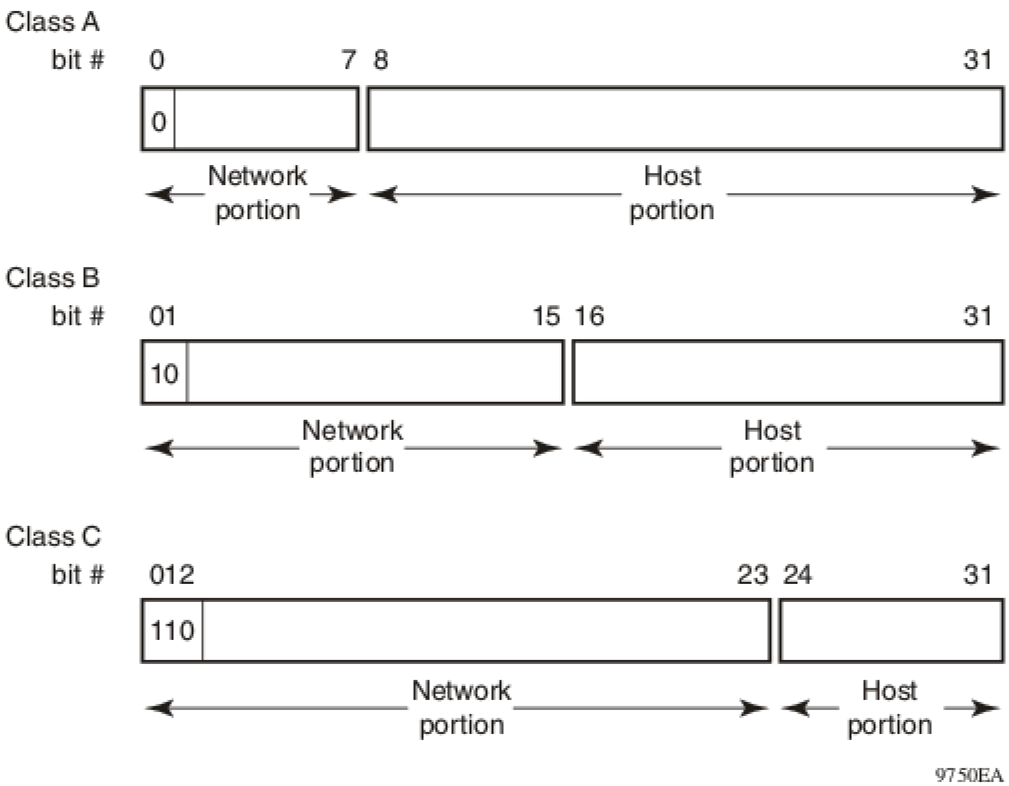

Figure 1: Network and host boundaries in IP address classes

Subnet addressing

Subnetworks (or subnets) are an extension of the IP addressing scheme. With subnets, organizations can use one IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

A subnet address is created by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is found in the first octet of the host portion (10). A subnet mask is applied to the IP address and identifies the network and host portions of the address.

The following table illustrates how subnet masks used with Class B and Class C addresses can create differing numbers of subnets and hosts. This example shows the use of the zero subnet, which is permitted on an Avaya Ethernet Routing Switch 4000 Series.

Table 3: Subnet masks for Class B and Class C IP addresses

Number of bits	Subnet Mask	Number of Subnets (Recommended)	Number of Hosts per Subnet
Class B			
2	255.255.192.0	2	16 382
3	255.255.224.0	6	8190
4	255.255.240.0	14	4094
5	255.255.248.0	30	2046
6	255.255.252.0	62	1022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1022	62
11	255.255.255.224	2046	30
12	255.255.255.240	4094	14
13	255.255.255.248	8190	6
14	255.255.255.252	16 382	2
Class C			
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

Variable-length subnet masking (VLSM) is the ability to divide an intranet into pieces that match network requirements. Routing is based on the longest subnet mask or network that matches.

IP routing

To configure IP routing on the Avaya Ethernet Routing Switch 4000 Series, you must create virtual router interfaces by assigning an IP address to a virtual local area network (VLAN). The following sections provide more details about IP routing functionality.

For a more detailed description about VLANs and their use, see *Configuring VLANs, Spanning Tree, and Multi-Link Trunking on Avaya Ethernet Routing Switch 4000 Series*, NN47205-501.

IP routing using VLANs

The Avaya Ethernet Routing Switch 4000 Series supports wire-speed IP routing between VLANs. To create a virtual router interface for a specified VLAN, you must associate an IP address with the VLAN.

The virtual router interface is not associated with any specific port. The VLAN IP address can be reached through any of the ports in the VLAN. The assigned IP address also serves as the gateway through which packets are routed out of that VLAN. Routed traffic can be forwarded to another VLAN within the switch or stack.

When the Avaya Ethernet Routing Switch 4000 Series is routing IP traffic between different VLANs, the switch is considered to be running in Layer 3 mode; otherwise, the switch runs in Layer 2 mode. When you assign an IP address to a Layer 2 VLAN, the VLAN becomes a routable Layer 3 VLAN. You can assign a single and unique IP address to each VLAN.

You can configure the global status of IP routing to be enabled or disabled on the Avaya Ethernet Routing Switch 4000 Series. By default, IP routing is disabled.

In this release, the Avaya Ethernet Routing Switch 4000 Series supports local routes and static routes. With local routing, the switch automatically creates routes to each of the local Layer 3 VLAN interfaces. With static routing, you must manually enter the routes to the destination IP addresses.

Local routes

With routing globally enabled, if you assign an IP address to a VLAN, IP routing is enabled for that VLAN. In addition, for each IP address assigned to a VLAN interface, the Ethernet Routing Switch adds a directly connected or local route to its routing table based on the IP address/mask assigned.

Local routing example

The following figure shows how the Ethernet Routing Switch can route between Layer 3 VLANs. In this example, the Ethernet Routing Switch has two VLANs configured. IP Routing is enabled globally on the switch and on the VLANs, each of which has an assigned IP address.

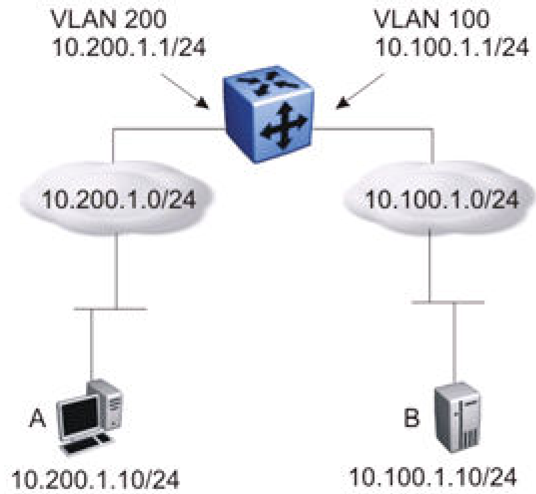

Figure 2: Local routes example

IP address 10.100.1.1/24 is assigned to VLAN 100, and IP address 10.200.1.1/24 is assigned to VLAN 200. As IP Routing is enabled, two local routes become active on the Avaya Ethernet Routing Switch as described in the following table.

	Network	Net-mask	Next-hop	Type
1	10.100.1.0	255.255.255.0	10.100.1.1	LOCAL
2	10.200.1.0	255.255.255.0	10.200.1.1	LOCAL

At this stage, both hosts A (10.200.1.10) and B (10.100.1.10) are reachable from the Ethernet Routing Switch. However, to achieve Layer 3 connectivity between A and B, additional configuration is required. Host A must know how to reach network 10.100.1.0/24, and host B must know how to reach network 10.200.1.0/24.

On host A, you must configure a route to network 10.100.1.0/24 through 10.200.1.1, or configure 10.200.1.1 as the default gateway for the host.

On host B, you must configure a route to network 10.200.1.0/24 through 10.100.1.1, or configure 10.100.1.1 as the default gateway for the host.

With these routes configured, the Ethernet Routing Switch can perform inter-VLAN routing, and packets can flow between hosts A and B.

Static routes

After you create routable VLANs through IP address assignment, you can create static routes. With static routes, you can manually create specific routes to destination IP addresses.

Non-local static routes

After you create routable VLANs through IP address assignment, you can create static routes. With static routes, you can manually create specific routes to destination IP addresses. Local routes have a next-hop that is on a directly connected network, while non-local routes have a next-hop that is not on a directly connected network. Non-local static routes are useful in situations where there are multiple paths to a network and the number of static routes can be reduced by using only one route with a remote gateway.

Static routes are not easily scalable. Thus, in a large or growing network this type of route management may not be optimal. Also, static routes do not have the capacity to determine the failure of paths. Thus, a router can still attempt to use a path after it has failed.

Static routes

After you create routable VLANs though IP address assignment, you can create static routes. With static routes, you can manually create specific routes to a destination IP address. In this release, the Ethernet Routing Switch supports local static routes only. For a route to become active on the switch, the next-hop IP address for the route must be on a directly connected network.

Static routes are not easily scalable. Thus, in a large or growing network, this type of route management may not be optimal.

Static routing example

The following figure shows an example of static routing on the Ethernet Routing Switch.

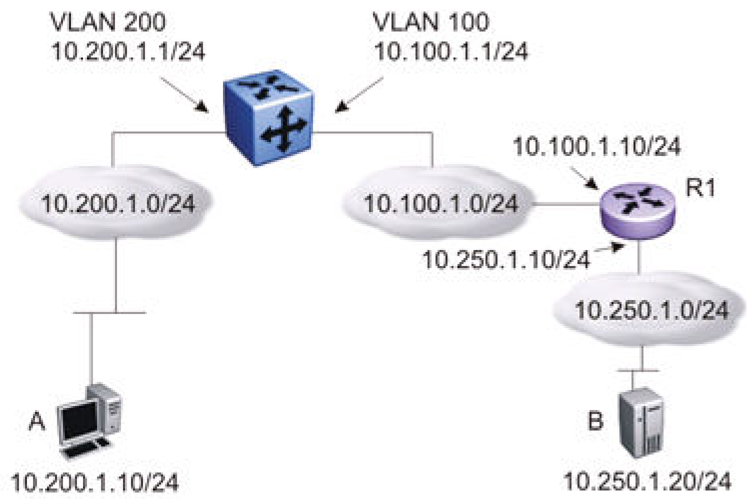

Figure 3: Static routes

In this example, two Layer 3 devices are used to create a physical link between hosts A and B. This network contains an Ethernet Routing Switch and another Layer 3 router, R1.

In this setup, the local route configuration from Local routing example on page 26 still applies. However, in this case, network 10.100.1.0/24 stands in between networks 10.200.1.0/24 and 10.250.1.0/24. To achieve end-to-end connectivity, router R1 must know how to reach network 10.200.1.0/24, and the Ethernet Routing Switch must know how to reach network 10.250.1.0/24. On the Ethernet Routing Switch, you can accomplish this using static routing. With static routing, you can configure a route to network 10.250.1.0/24 through 10.100.1.10. In this case, the following routes are active on the Ethernet Routing Switch.

	Network	Net-mask	Next-hop	Type
1	10.100.1.0	255.255.255.0	10.100.1.1	LOCAL
2	10.200.1.0	255.255.255.0	10.200.1.1	LOCAL
3	10.250.1.0	255.255.255.0	10.100.1.10	STATIC

To obtain Layer 3 connectivity between the hosts, additional routes are required. Host A requires a route to 10.250.1.0/24 using 10.200.1.1 as the next hop, or with 10.200.1.1 as the default gateway. Host B requires a route to 10.200.1.0/24 using 10.250.1.10 as the next hop, or with 10.250.1.10 as the default gateway.

The configuration for router R1 to reach network 10.200.1.0/24 is dependent on the type of router used.

Default routes

Default routes specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This static default route is a route to the network address 0.0.0.0 as defined by the Institute of Electrical and Electronics Engineers (IEEE) Request for Comment (RFC) 1812 standard.

The Ethernet Routing Switch uses the default route 0.0.0.0/0.0.0.0 for all Layer 3 traffic that does not match a specific route. This traffic is forwarded to the next-hop IP address specified in the default route.

Route scaling

The Avaya Ethernet Routing Switch 4000 Series supports a maximum of 256 local routes and up to 256 static routes, including the default route (Destination = 0.0.0.0, Mask = 0.0.0.0). The partitioning of the route table can be altered using the Dynamic Routing Table Allocation feature.

Dynamic Routing Table Allocation

Dynamic Routing Table Allocation increases the flexibility of the Avaya Ethernet Routing Switch 4000 to support different combinations of static and dynamic routing protocols like Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). In combination with Route Policies, Dynamic Routing Table Allocation can allow you to efficiently utilize switch resources. With Dynamic Route Table Allocation, you can manually partition the maximum 512 IPv4 route entries between the various routing protocols supported on the switch. When you use Dynamic Route Table Allocation, you can reserve from 2 to 256 entries for local routes and 0 to 256 entries for static routes (including non-local static routes), with the remainder of the 512 IPv4 route entries (to a maximum of 510) available for RIP and OSPF.

The default configuration supports the availability of 64 local routing instances, 32 static routes, and 416 dynamic routes for RIP and OSPF. Additional dynamic route entries can be allocated when IP routing is disabled by decreasing the number of local and static routes. The number of local routes cannot be reduced below the number of Layer 3 VLANs or routing instances configured on the switch.

If you decrease the combined maximum totals of local and static routes, the total number of available dynamic routes is increased. You can increase the total number of available dynamic routes whether IP routing is enabled or disabled.

If you increase the combined maximum totals of local and static routes, the total number of available dynamic routes is decreased. You can decrease the total number of available dynamic routes only when IP routing is disabled. Therefore, the following considerations apply when you change the maximum values for local and static routes with IP routing enabled:

- If you increase the maximum number of local routes (n), you must decrease the maximum number of static routes by the same amount (n).
- If you increase the maximum number of static routes (n), you must decrease the maximum number of local routes by the same amount (n).

If IP routing is disabled, you can change the maximum totals of local and static routes without limitations.

Management VLAN

With IP routing enabled on the switch or stack, you can use any of the virtual router IP addresses for device management over IP. Any routable Layer 3 VLAN can carry the management traffic for the switch, including Telnet, Web, Simple Network Management Protocol (SNMP), BootP, and Trivial File Transfer Protocol (TFTP). Without routing enabled, the management VLAN is reachable only through the switch or stack IP address, and only through ports that are members of the management VLAN. The management VLAN always exists on the switch and cannot be removed.

When routing is enabled on the Avaya Ethernet Routing Switch 4000 Series switches, the management VLAN behaves similar to other routable VLANs. The IP address is reachable through any virtual router interface, as long as a route is available.

Management route

On the Ethernet Routing Switch, you can configure a management route from the Management VLAN to a particular subnet. The management route is a static route that allows incoming management connections from the remote network to the management VLAN.

The management route transports traffic between the specified destination network and the Management VLAN only. It does not carry inter-VLAN routed traffic from the other Layer 3 VLANs to the destination network. This provides a management path to the router that is inaccessible from the other Layer 3 VLANs. While you can access the management VLAN from all static routes, other static routes cannot route traffic to the management route.

To allow connectivity through a management route, you must enable IP routing globally and on the management VLAN interface.

The following figure shows an example of a management route allowing access to the management VLAN interface.

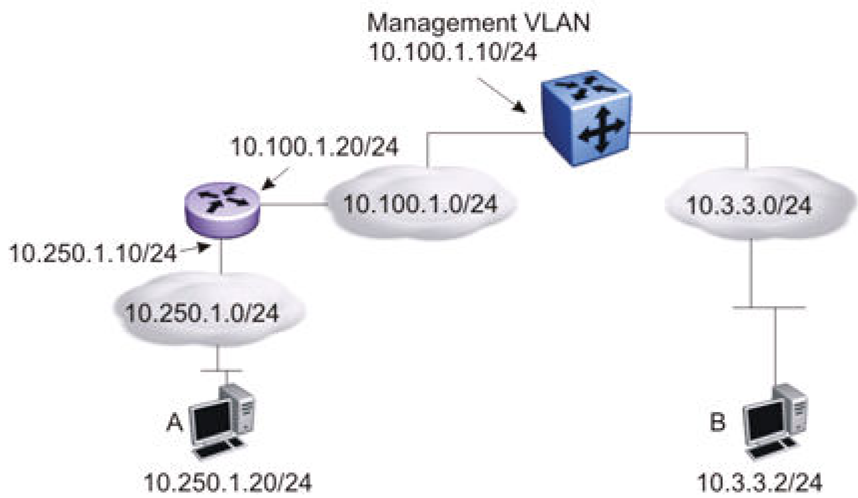

Figure 4: Management route

As network 10.250.1.0/24 is not directly connected to the Ethernet Routing Switch, to achieve connectivity from host 10.250.1.20 to the management VLAN, the Ethernet Routing Switch must know how to reach network 10.250.1.0/24. On the Ethernet Routing Switch, you can configure a management route to network 10.250.1.0/24 through 10.100.1.20. In this case, the following management route is active on the Ethernet Routing Switch.

	Network	Net-mask	Next-hop	Type
1	10.250.1.0	255.255.255.0	10.100.1.20	MANAGEMENT

With this configured route, host A at 10.250.1.20 can perform management operations on the Ethernet Routing Switch. To do so, Host A also requires a route to 10.100.1.0/24 using 10.250.1.10 as the next hop, or with 10.250.1.10 as the default gateway.

If a Layer 3 VLAN is also configured for network 10.3.3.0/24, this provides a local route that host B at 10.3.3.2 can use to access the switch. However, host B cannot communicate with host A, as the route to network 10.250.1.0/24 is a management route only. To provide connectivity between the two hosts, you must configure a static route to 10.250.1.0/24.

Brouter port

You can configure brouter ports for the ERS 4000 switches. A brouter port is a single-port VLAN that can route IP packets as well as bridge all non-routable traffic. The difference between a brouter port and a standard IP protocol-based VLAN configured for routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for non-routable traffic and route IP traffic, thereby removing potential interruptions caused by Spanning Tree Protocol recalculations in routed traffic. A brouter port is a one-port VLAN; each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

When you create a brouter port the system performs the following actions on the switch:

- A port-based VLAN is created.
- The brouter port is added to the new port-based VLAN.
- The PVID of the brouter port is changed to the VLAN ID of the new VLAN.
- The brouter VLAN is added to a new STP group which is hidden to the user. The port is in forwarding state all the time in this new STP group (the spanning-tree protocol does not apply for this group). The port is in forwarding state from the beginning without setting the STP participation to disabled in the default STP group.
- An IP address is assigned to the brouter VLAN.

Related routing features

The following sections describe features that are related to and dependent on the IP routing functionality.

DHCP relay

Dynamic Host Configuration Protocol (DHCP) is a mechanism to assign network IP addresses on a dynamic basis to clients who request an address. DHCP is an extension of the Bootstrap protocol (BootP). BootP/DHCP clients (workstations) generally use User Datagram Protocol (UDP) broadcasts to determine their IP addresses and configuration information. If such a host is on a VLAN that does not include a DHCP server, the UDP broadcasts are by default not forwarded to servers located on different VLANs.

The Avaya Ethernet Routing Switch 4000 Series can resolve this issue using DHCP relay, which forwards the DHCP broadcasts to the IP address of the DHCP server. Network managers prefer to configure a small number of DHCP servers in a central location to lower administrative overhead. Routers must support DHCP relay so that hosts can access configuration information from servers several router hops away.

With DHCP relay enabled, the switch can relay client requests to DHCP servers on different Layer 3 VLANs or in remote networks. It also relays server replies back to the clients.

To relay DHCP messages, you must create two Layer 3 VLANs: one connected to the client and the other providing a path to the DHCP server. You can enable DHCP relay on a per-VLAN basis.

The following figure shows a DHCP relay example, with an end station connected to subnet 1, corresponding to VLAN 1. The Avaya Ethernet Routing Switch 4000 Series connects two subnets by means of the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255), with the DHCP relay function enabled, the Ethernet Routing Switch forwards the DHCP request to the host address of the DHCP server on VLAN 2.

Figure 5: DHCP relay operation

Forwarding DHCP packets

In the following figure, the DHCP relay agent address is 10.10.1.254. To configure the Avaya Ethernet Routing Switch 4000 Series to forward DHCP packets from the end station to the server, use 10.10.2.1 as the server address.

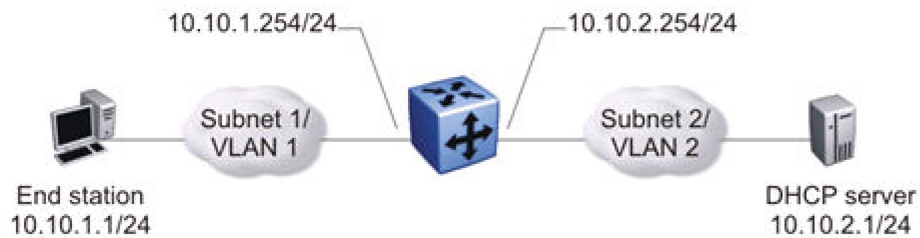

Figure 6: Forwarding DHCP packets

All BootP and DHCP broadcast packets that appear on the VLAN 1 router interface (10.10.1.254) are then forwarded to the DHCP server. In this case, the DHCP packets are forwarded as unicast to the DHCP server IP address.

Multiple DHCP servers

Most enterprise networks use multiple DHCP servers for fault tolerance. The Avaya Ethernet Routing Switch 4000 Series can forward DHCP requests to multiple servers. You can configure up to 256 servers to receive copies of the forwarded DHCP messages.

To configure DHCP client requests to be forwarded to multiple different server IP addresses, specify the client VLAN as the DHCP relay agent for each of the destination server IP addresses.

In the following figure, two DHCP servers are located on two different VLANs. To configure the Avaya Ethernet Routing Switch 4000 Series to forward copies of the DHCP packets from the end station to both servers, specify the IP address of VLAN 1 (10.10.1.254) as the DHCP relay agent address and associate this relay agent with each of the DHCP server addresses, 10.10.2.1 and 10.10.3.1.

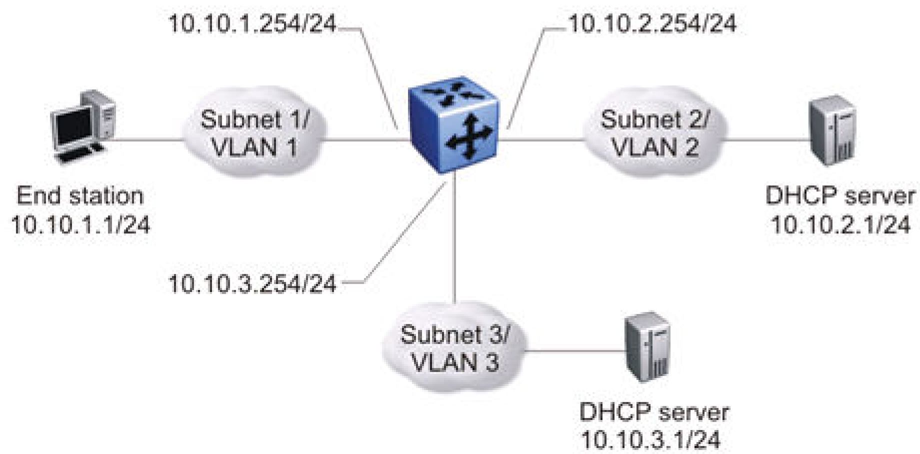

Figure 7: Multiple DHCP servers

Differences between DHCP and BootP

With DHCP relay, the Avaya Ethernet Routing Switch 4000 Series supports the relay of DHCP and the Bootstrap protocol (BootP). The following differences between DHCP and BootP are specified in RFC 2131:

- BootP enables the retrieval of an American Standard Code for Information Interchange (ASCII) configuration file name and configuration server address.
- A properly configured BootP server enables the switch to automatically learn its assigned IP address, subnet mask, and the IP address of the default router (default gateway).
- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters they need to operate.

DHCP uses the BootP message format defined in RFC 951. The remainder of the options field consists of a list of tagged parameters that are called options(RFC 2131).

DHCP Option 82

With DHCP Option 82, the switch can optionally add information about the client port when relaying the DHCP request to the DHCP server. This information from the switch can be used to identify the location of the device in the network. DHCP Option 82 function is added by the switch at the edge of a network.

When a VLAN is operating in Layer 2 mode, DHCP Snooping must be enabled for DHCP Option 82 to function. When a VLAN is operating in Layer 3 (IP Routing) mode, the DHCP

Option 82 function requires that DHCP Relay is appropriately configured. To use DHCP Option 82 with DHCP relay, you must enable DHCP relay globally on the switch and client VLANs.

For information about DHCP Option 82 with DHCP snooping, see *Configuring Security on Avaya Ethernet Routing Switch 4000 Series*, NN47205-505.

DHCP Relay Packet Size

In accordance with RFC3046, the Avaya Ethernet Routing Switch 4000 provides the capability to specify the maximum frame size the DHCP relay agent forwards to the DHCP server. The switch implementation permits configuration of the maximum DHCP packet size to 1536 bytes, the default maximum size is 576 bytes. If the DHCP packet exceeds the maximum configured size, the DHCP Option 82 information is not appended to the message.

UDP broadcast forwarding

By default, User Datagram Protocol (UDP) broadcast frames received on one VLAN are not routed to another VLAN. To allow UDP broadcasts to reach a remote server, the Ethernet Routing Switch supports UDP broadcast forwarding, which forwards the broadcasts to the server through a Layer 3 VLAN interface.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address. The packet is sent as a unicast packet to the server.

When a UDP broadcast is received on a router interface, it must meet the following criteria to be considered for forwarding:

- It must be a MAC-level broadcast.
- It must be an IP-limited broadcast.
- It must be for a configured UDP protocol.
- It must have a time-to-live (TTL) value of at least 2.

For each ingress interface and protocol, the UDP broadcast packets are forwarded only to a unicast host address (for example, to the unicast IP address of the server).

When the UDP forwarding feature is enabled, a filter is installed that compares the UDP destination port of all packets against all the configured UDP forwarding entries. If a match occurs, the destination IP of the incoming packet is checked for consistency with the user-configured broadcast mask value for this source VLAN. If these conditions are met, the TTL field from the incoming packet is overwritten with the user-configured TTL value, the destination IP of the packet is overwritten with the configured destination IP, and the packet is routed to the destination as a unicast frame.

Important:

UDP broadcast forwarding shares resources with the Quality of Service (QoS) feature. When UDP forwarding is enabled, the switch dynamically assigns the highest available precedence value to the UDP forwarding feature. To display the assigned precedence after you enable UDP forwarding, enter the **show qos diag** command.

For further information on QoS policies, see *Configuring Quality of Service on Avaya Ethernet Routing Switch 4000 Series*, NN47205-504.

UDP forwarding example

Figure 8: UDP forwarding example on page 37 shows an example of UDP broadcast forwarding. In this case, if host A (10.200.1.10) needs a certain service (for example, a custom application that listens on UDP port 12345), it transmits a UDP broadcast frame. By default, the Ethernet Routing Switch does not forward this frame to VLAN 100, and because server B (10.100.1.10) is not on VLAN 200, the host cannot access that service.

With UDP broadcast forwarding enabled, the host can access the service. In this case, you must list port 12345 as a valid forwarding port, and specify VLAN 200 as the source VLAN.

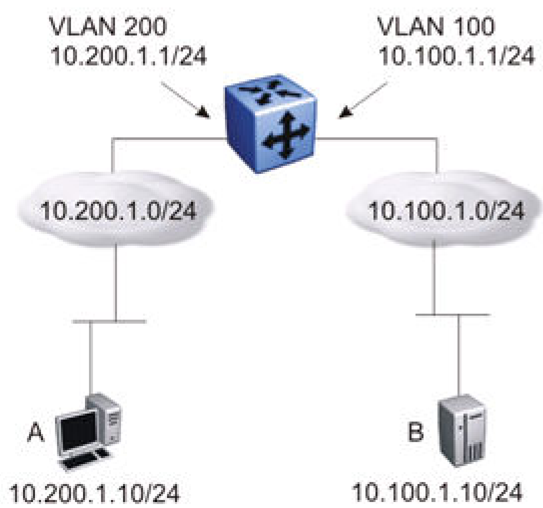

Figure 8: UDP forwarding example

When the switch receives an incoming packet on VLAN 200 that matches the configured UDP destination port (12345), and the destination IP is consistent with the broadcast mask value for the VLAN; then the switch applies the new destination IP (here, 10.100.1.10) to the packet and routes it to the destination as a unicast frame.

Directed broadcasts

With the directed broadcasts feature enabled, the Ethernet Routing Switch can determine if an incoming unicast frame is a directed broadcast for one of its interfaces. If so, the switch forwards the datagram onto the appropriate network using a link-layer broadcast.

With IP directed broadcasting enabled on a VLAN, the Ethernet Routing Switch forwards direct broadcast packets in the following two ways:

- through a connected VLAN subnet to another connected VLAN subnet
- through a remote VLAN subnet to the connected VLAN subnet

By default, this feature is disabled.

ARP

The Address Resolution Protocol (ARP) allows the Ethernet Routing Switch to dynamically learn Layer 2 Media Access Control (MAC) addresses, and to build a table with corresponding Layer 3 IP addresses.

Network stations using the IP protocol need both a physical (MAC) address and an IP address to transmit a packet. If a network station knows only the IP address of a network host, ARP enables the network station to determine the physical address of the network host and bind the 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts.

If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the physical address of the host as follows:

1. The network station broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.
2. All network hosts receive the broadcast message.
3. Only the specified host responds with its hardware address.
4. The network station then maps the host IP address to its physical address and saves the results in an address resolution table for future use.
5. The network station ARP table displays the association of the known MAC addresses to IP addresses.

The lifetime for the learned MAC addresses is a configurable parameter. The switch executes ARP lookups when this timer expires.

The default timeout value for ARP entries is 6 hours.

Static ARP

In addition to the dynamic ARP mechanism, the Ethernet Routing Switch supports a static mechanism that allows for static ARP entries to be added. With Static ARP, you can manually associate a device MAC address to an IP address. You can add and delete individual static ARP entries on the switch.

Proxy ARP

Proxy ARP allows the Ethernet Routing Switch to respond to an ARP request from a locally attached host that is intended for a remote destination. It does so by sending an ARP response back to the local host with the MAC address of the switch interface that is connected to the host subnet. The reply is generated only if the switch has an active route to the destination network.

With Proxy ARP enabled, the connected host can reach remote subnets without the need to configure default gateways.

The following figure is an example of proxy ARP operation. In this example, host B wants to send traffic to host C, so host B sends an ARP request for host C. However, the Avaya Ethernet Routing Switch 4000 Series is between the two hosts, so the ARP message does not reach host C. To enable communication between the two hosts, the Avaya Ethernet Routing Switch 4000 Series intercepts the message and responds to the ARP request with the IP address of host C but with the MAC address of the switch itself. Host B then updates its ARP table with the received information.

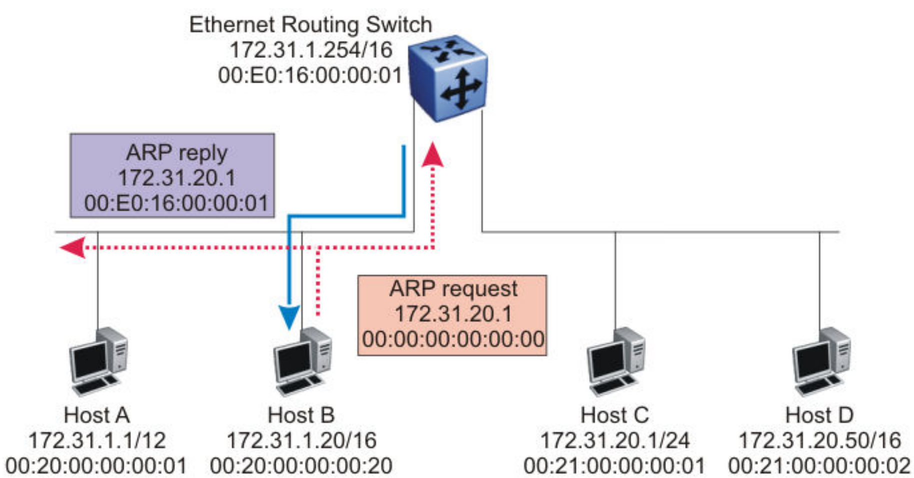

Figure 9: Proxy ARP Operation

Avaya recommends Proxy ARP as a temporary fix only, for example, if you are gradually moving hosts from one addressing scheme to another and you still want to maintain connectivity between the disparately-addressed devices. You do not want Proxy ARP running as a general rule because it causes hosts to generate ARP messages for every address that they want to reach on the Internet.

IP blocking for stacks

IP blocking is a Layer 3 feature of the Avaya Ethernet Routing Switch 4000 Series that provides safeguards for a stack where Layer 3 VLANs have port members across multiple stack units.

IP Blocking is used whenever a unit leaves a stack or is rebooting inside the context of a stack. Depending on the setting in use, Layer 3 functionality is either continued or blocked by this feature.

You can set the IP Blocking mode on the base unit to either none or full.

When IP blocking is set to full, if any units leave the stack, those units run in Layer 2 mode. No Layer 3 settings remain on the units.

When IP blocking is set to none, if any units leave the stack, the Layer 3 configurations applied to the stack are still applied on the individual units.

In a stack environment of 2 units, Avaya recommends that you use IP blocking mode none. In this case, you can expect the following functional characteristics:

- If either the stack base unit or nonbase unit becomes nonoperational, Layer 3 functionality continues to run on the remaining unit.

A disadvantage of this configuration is that if the nonoperational unit does not rejoin the stack, address duplication occurs.

In stack environments of more than 2 units, Avaya recommends that you use IP blocking mode full. In this case, you can expect the following functional characteristics:

- If the stack base unit becomes nonoperational, the following occurs:
    - The temporary base unit takes over base unit duties.
    - The temporary base unit takes over responsibility to manage Layer 3 functionality in the stack. When this occurs, the system updates the MAC addresses associated with each routing interface to be offset from the temporary base unit MAC address (rather than the base unit MAC address). During this period, some minor disruption may occur to routing traffic until end stations update their ARP cache with the new router MAC addresses. The Avaya Ethernet Routing Switch 4000 Series sends out gratuitous ARP messages on each routed VLAN for 5 minutes at 15 second intervals to facilitate quick failover in this instance.
    - If the nonoperational base unit does not rejoin the stack, no Layer 3 functionality runs on the unit.
- If a stack nonbase unit becomes nonoperational, the following occurs:
    - The stack continues to run normally with the base unit controlling Layer 3 functionality.
    - If the nonoperational nonbase unit does not rejoin the stack, no Layer 3 functionality runs on the unit.

By default, the IP blocking mode is none (disabled).

Open Shortest Path First (OSPF) protocol

Open Shortest Path First (OSPF) is a classless Interior Gateway Protocol (IGP) that distributes routing information between routers belonging to a single autonomous system (AS). An OSPF AS is generally defined as a group of routers in a network that run OSPF and that operate under the same administration. Intended for use in large networks, OSPF is a link-state protocol that supports variable length subnet masking (VLSM) and tagging of externally-derived routing information.

Important:

The Avaya Ethernet Routing Switch 4000 Series implementation of OSPF supports broadcast and passive interfaces. The NBMA type interfaces are not supported.

Overview

In an OSPF network, each router maintains a link-state database that describes the topology of the autonomous system (AS). The database contains the local state for each router in the AS, including usable interfaces and reachable neighbors. Each router periodically checks for changes in its local state and shares detected changes by flooding link-state advertisements (LSA) throughout the AS. Routers synchronize their topological databases based on the sharing of information from LSAs.

From the topological database, each router constructs a shortest-path tree, with itself as the root. The shortest-path tree gives the optimal route to each destination in the AS. Routing information from outside the AS appears on the tree as leaves.

In large networks, OSPF offers the following benefits:

- Provides support for different routing authentication methods to guard against passive attacks
- Recalculates routes quickly during the network topology change
- Generates a minimum of routing protocol traffic
- Provides support for equal-cost multipath routing. If several equal-cost routes to a destination exist, it distributes the traffic equally among them.
- Offers scalable routing domain because it does not use hop count in its calculation
- Allows you to import external routes (RIP, BGP) into OSPF domain
- Allows large network to be partitioned into smaller and contiguous areas
- Provides mechanism for aggregation routes between areas that help in reducing routing table size, network bandwidth, and CPU utilization
- Uses IP multicast to discover neighbors and send link-state updates

OSPF routes IP traffic based on the destination IP address, subnet mask, and IP TOS.

Autonomous system and areas

In large OSPF networks with many routers and networks, the link-state database (LSDB) and routing table on each router can become excessively large. Large route tables and LSDBs consume memory. In addition, the processing of additional LSAs puts added strain on the CPU to make forwarding decisions. To reduce these undesired effects, an OSPF network can be divided into subdomains called areas. Each area comprises a number of OSPF routers that have the same area ID. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic compared to treating the entire AS as a single link-state domain.

When a network is divided into multiple areas, each router within an area maintains an LSDB only for the area to which it belongs. Each area is identified by a unique 32-bit area ID, expressed in IP address format (x.x.x.x). Area 0.0.0.0 is known as the backbone area and distributes routing information to all other areas.

Within the AS, packets are routed based on their source and destination addresses. If the source and destination of a packet reside in the same area, intra-area routing is used. Intra-area routing protects the area from bad routing information because no routing information obtained from outside the area can be used.

If the source and destination of a packet reside in different areas, inter-area routing is used. Inter-area routing must pass through the backbone area.

ABR

A router attached to two or more areas inside an OSPF network is identified as an Area Border Router (ABR). Each ABR maintains a separate topological database for each connected area. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information from one area to another. When the AS is divided into multiple areas, each nonbackbone area must be attached to the backbone area through an (ABR).

For routers that are internal to an area (identified as internal routers), the impact of a topology change is localized to the area in which it occurs. However, ABRs must maintain an LSDB for each area to which they belong. ABRs advertise changes in topology from one area to another by advertising summary LSAs.

Backbone area

The backbone area connects nonbackbone areas to each other. Traffic forwarded from one area to another must travel through the backbone. The backbone topology dictates the paths used between areas. The topology of the backbone area is invisible to other areas and the backbone has no knowledge of the topology of nonbackbone areas.

The area ID 0.0.0.0 is reserved for the backbone area.

Area border routers (ABR) cannot learn OSPF routes unless they have a connection to the backbone. Inter-area paths are selected by examining the routing table summaries for each connected ABR.

In inter-area routing, a packet travels along three contiguous paths:

1. First, the packet follows an intra-area path from the source to an ABR, which provides the link to the backbone.
2. From the source ABR, the packet travels through the backbone toward the destination area ABR.
3. At the destination area ABR, the packet takes another intra-area path to the destination.

The following figure shows an OSPF AS divide into three areas: a backbone area, a stub area, and a not-so-stubby area (NSSA). (Stub areas and NSSAs are described in subsequent sections.)

The figure also shows ABRs connecting the areas to one another and Autonomous System Border Routers (ASBR) connecting two areas to external networks. ASBRs redistribute external static or RIP routes into the OSPF network.

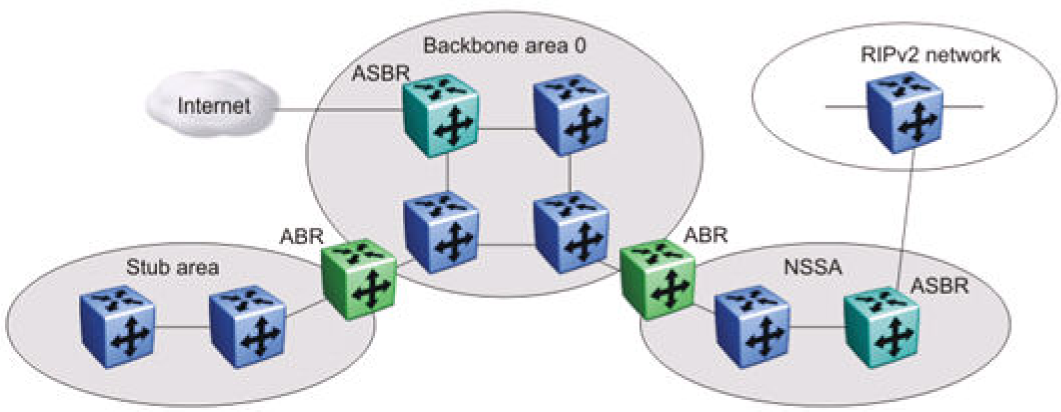

Figure 10: OSPF network

OSPF neighbors

In an OSPF broadcast network, any two routers that have an interface to the same network are neighbors. OSPF routers use the Hello Protocol to dynamically discover and maintain neighbor relationships.

Periodically, OSPF routers send Hello packets over all interfaces to the AllSPFRouters multicast address. These Hello packets include the following information:

- router priority
- router Hello Timer and Dead Timer values

- list of routers that sent the router Hello packets on this interface
- router choice for designated router (DR) and backup designated router (BDR)

Bidirectional communication is determined when a router discovers itself listed in its neighbor Hello packet.

Designated routers

To form an adjacency, two OSPF routers perform a database exchange process to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent.

To limit the amount of routing protocol traffic, OSPF routers use the Hello Protocol to elect a designated router (DR) and a backup designated router (BDR) on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information (which on a large network can mean significant routing protocol traffic), all routers on the network form adjacencies with the DR and the BDR only, and send link-state information only to them. The DR redistributes this information to every other adjacent router.

The BDR receives link-state information from all routers on the network and listens for acknowledgements. If the DR fails, the BDR can transition quickly to the role of DR because its routing tables are up to date.

OSPF Operation

On broadcast multiaccess networks, the sequence of processes governed by OSPF is as follows:

1. When a router starts, it initializes the OSPF data structures and then waits for indications from lower-level protocols that the router interfaces are functional.
2. The router dynamically detects neighbors by sending and receiving Hello packets to the AllSPFRouters multicast address.
3. Using the Hello Protocol, a designated router (DR) and backup designated router (BDR) are elected for the network.
4. Each router forms an adjacency and exchanges database information only with the DR and the BDR.
5. The DR floods LSAs containing information about each router and its neighbors throughout the area to ensure that all routers in the area have an identical topological database.

6. From this database each router uses the OSPF routing algorithm (Dijkstra's algorithm) to calculate a shortest-path tree, with itself as root. This shortest-path tree in turn yields a routing table for the protocol.
7. After the network has converged, each OSPF router continues to periodically flood Hellos to maintain neighbor relationships. And at longer intervals, LSAs are retransmitted throughout the area. In addition, routers forwards LSAs to the DR if they detect a change in the state of a router or a link (that is, up or down). Upon receipt of an LSA, the DR can then flood the update to all routers in the area, enabling quick detection of dead routers on the network.

OSPF route advertisements

A destination in an OSPF route advertisement is expressed as an IP address and a variable-length mask. Together, the address and the mask indicate the range of destinations to which the advertisement applies.

Because OSPF can specify a range of networks, it can send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 128.185.0.0 with a mask of 255.255.0.0 describes a single route to destinations 128.185.0.0 to 128.185.255.255.

Router types

As mentioned in preceding sections, routers in an OSPF network can have various roles depending on how you configure them. The following table describes the router types you can configure in an OSPF network.

Table 4: Router types in an OSPF network

Router type	Description
AS boundary router (ASBR)	A router attached at the edge of an OSPF network is called an ASBR. Any router that distributes static routes or RIP routes into OSPF is considered an ASBR. The ASBR forwards external routes into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area border router (ABR)	A router attached to two or more areas inside an OSPF network is considered an ABR. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information.
Internal router (IR)	A router that has interfaces only within a single area inside an OSPF network is considered an IR. Unlike ABRs, IRs have topological information only about the area in which they are contained.

Router type	Description
Designated router (DR)	In a broadcast network, a single router is elected to be the DR for that network. A DR ensures that all routers on the network are synchronized and advertises the network to the rest of the AS.
Backup designated router (BDR)	A BDR is elected in addition to the DR and, if the DR fails, can assume the DR role quickly.

LSA types

After the network has converged, OSPF does not require each router to keep sending its entire LSDB to its neighbors. Instead, each OSPF router floods only link-state change information in the form of LSAs throughout the area or AS. LSAs typically contain information about the router and its neighbors and are generated periodically to ensure connectivity or are generated by a change in state of the router or a link (that is, up or down).

The following table displays the seven LSA types exchanged between OSPF routers.

Table 5: OSPF LSA types

LSA type	LSA name	Description	Area of distribution
1	Router LSA	Type 1 LSAs are originated by every router to describe their set of active interfaces and neighboring routers. Type 1 LSAs are flooded only within the area. A backbone router can flood router link advertisements within the backbone area.	Only within the same area
2	Network LSA	Type 2 LSAs describe a network segment. In a broadcast network, the designated router (DR) originates network LSAs that list all routers on that LAN. Type 2 LSAs are flooded only within the area. A backbone DR can flood network links advertisements within the backbone area.	Only within the same area
3	Network-Summary LSA	Type 3 LSAs are originated by the area border router (ABR) to describe the networks that are reachable outside the area. An ABR attached to two areas generates a different network summary LSA for each area. ABRs also flood type 3 LSAs containing information about destinations within an area to the backbone area.	Passed between areas

LSA type	LSA name	Description	Area of distribution
4	ASBR-summary LSA	Type 4 LSAs are originated by the ABR to advertise the cost of the path to the closest ASBR from the router generating the advertisement.	Passed between areas
5	Autonomous System External [ASE] LSA	Type 5 LSAs are originated by the ASBR to describe the cost of the path to a destination outside the AS from the ASBR generating the advertisement. Type 5 LSAs are passed between areas. In stub and NSSA areas, type 5 LSA routes are replaced with a single default route.	Passed between areas
6	Group Membership LSA	Type 6 LSAs identify the location of multicast group members in multicast OSPF.	Passed between areas
7	NSSA External LSA	Type 7 LSAs are used in OSPF NSSAs to import external routes.	Translated between areas

Area types

OSPF supports multiple area types. The following sections describe the supported OSPF area types.

Stub area

As shown in the following figure, a stub area is configured at the edge of the OSPF routing domain and has only one ABR.

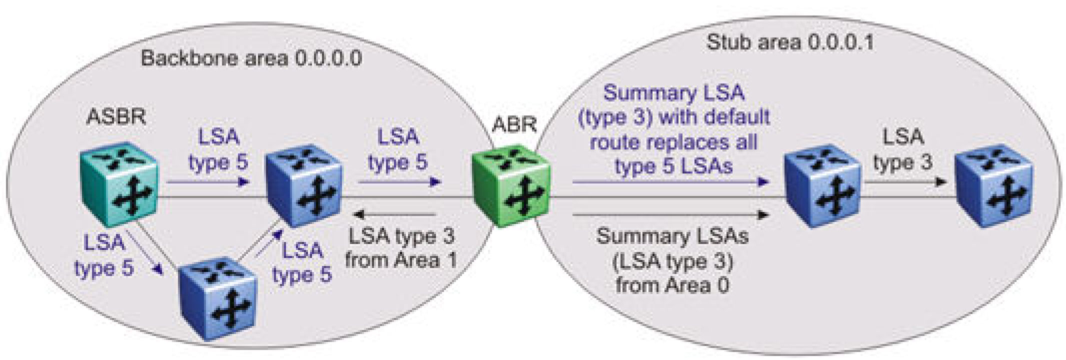

Figure 11: Stub area

The ABR does not flood AS External LSAs (type 5) into a stub area. Instead, the ABR uses Summary LSAs (type 3) to advertise a default route (0.0.0.0) into the stub area for all external routes. As stub areas do not receive advertisements for external routes from the ABR, the size of the link state database in the stub area is reduced.

For internal routers in the stub area, any destinations that do not match intra-area or inter-area routes are passed to the ABR for routing to the external destinations.

Because stub areas do not support type 5 ASE LSAs, they cannot support ASBRs.

Not so stubby area

Like a stub area, a not so stubby area (NSSA) is at the edge of an OSPF routing domain and it prevents the flooding of AS External LSAs into the NSSA by replacing them with a default route.

However, unlike a stub area, an NSSA can import small stub (non-OSPF) routing domains into OSPF. This allows the NSSA to import external routes, such as RIP routes, and advertise these routes throughout the network.

As shown in the following figure, a non-OSPF routing domain can connect to the NSSA to allow the external network to route traffic to the OSPF AS. One router in the NSSA must operate as an ASBR to provide a link to the non-OSPF domain.

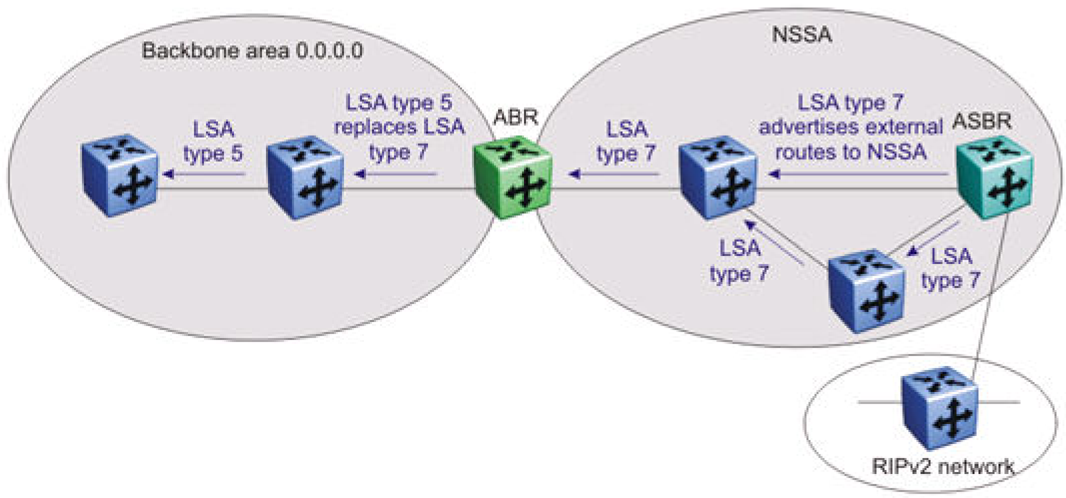

Figure 12: OSPF NSSA

If the non-OSPF network is a small network, and the attached non-OSPF router has a default route to the OSPF network, this provides sufficient routing for any destinations that are outside the non-OSPF network.

Within the NSSA, the NSSA ASBR advertises route information imported from the external network using type 7 LSAs (NSSA External LSAs).

To propagate the external routes to other areas, the NSSA ABR translates these type 7 LSAs into type 5 LSAs (AS External LSAs). The ABR can flood the type 5 LSAs to the other areas so that the rest of the OSPF domain can learn about the non-OSPF destinations.

You can also configure the ABR to prevent the flooding of the external routes to other areas. To support this additional control over external route advertisement, the type 7 LSAs provide an Options field containing an N/P-bit that notifies the ABR which external routes can be advertised to other areas. When the NSSA N/P-bit is set to true (the default setting), the ABR exports the external route. When the NSSA N/P-bit is not set, the ABR drops the external route.

To manipulate the N/P-bit value for specific routes, you must configure a route policy on the Avaya Ethernet Routing Switch 4000 Series.

Normal area

A normal area is an area that is neither a bacbone nor a stub area that sends and receives LSA types 1 through 5. As illustrated in the following figure, a normal area supports Area Border Routers (ABRs) and Autonomous System Border Routers (ASBRs).

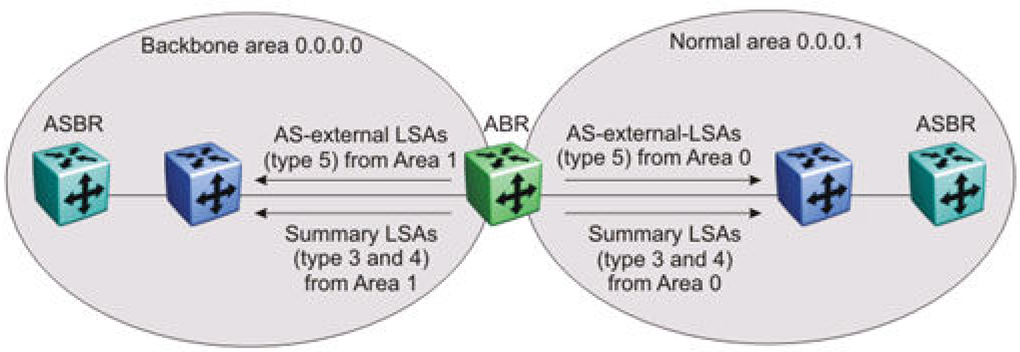

Figure 13: OSPF normal area

The Avaya Ethernet Routing Switch 4000 Series automatically becomes an ABR when it is connected to more than one area.

Area aggregation

OSPF maintains a table of area aggregation range configured for each area.

The area aggregation

- automatically places the OSPF routing interface into a specific area.
- advertise or suppress summary LSA for group of subnets to reduce the number of OSPF summary packets between areas, and to conserve router memory needed for link-state database.

The table maintains information in terms of area ID, LSA type (summary-link/nssa-extlink), and network address.

You can configure multiple area aggregate ranges for the same area, thus, OSPF summarizes addresses for many different set of address ranges. Avaya Ethernet Routing Switch 4000 Series implemetation of OSPF allows you to configure up to eight range for each area.

The following advertise modes are supported:

- Summarize ABR

  Sends only one summary LSA for all networks that fall within the range
- Suppress ABR

  Does not send any summary LSA for networks that fall within the range
- No Summarize ABR

  Sends summary LSAs for individual networks within the range

Advertise metric is the cost value that you want to advertise for the OSPF area range.

SPF calculation

The Avaya Ethernet Routing Switch 4000 Series uses the Dijkstra algorithm to calculate the shortest path. In this algorithm, the shortest path from a router to each known destination is calculated based on the cumulative cost required to reach that destination. This algorithm takes link-state database as input, and performs a separate calculation for each area the router belongs to. After completing the calculation, the router updates the routing table.

If there is a topology change, the SPF calculation is triggered automatically. You can also start it manually by setting a system parameter.

The following types of route calculations are required depending on the types of topology changes

- Intra-area route computation
- Inter-area route computation
- External route computation

The following events trigger recalculation of OSPF routes upon expiration of the configurable holddown timer:

- Update or new router-LSA and network-LSA
- Update or new summary-LSA
- New external-route-LSA
- Manual setting of SPF run flag

OSPF virtual link

The OSPF network can be partitioned into multiple areas. However, every non-backbone area must be connected to the backbone area through an ABR. If no physical connection to the backbone is available, you can create a virtual link.

A virtual link is established between two ABRs and is a logical connection to the backbone area through a non-backbone area called a transit area. Stub or NSSA areas cannot be transit areas.

In the following diagram, non-backbone ABR R4 establishes a virtual link with backbone ABR R1 across transit area 1.1.1.1. The virtual link connects area 2.2.2.2 to area 0.0.0.0.

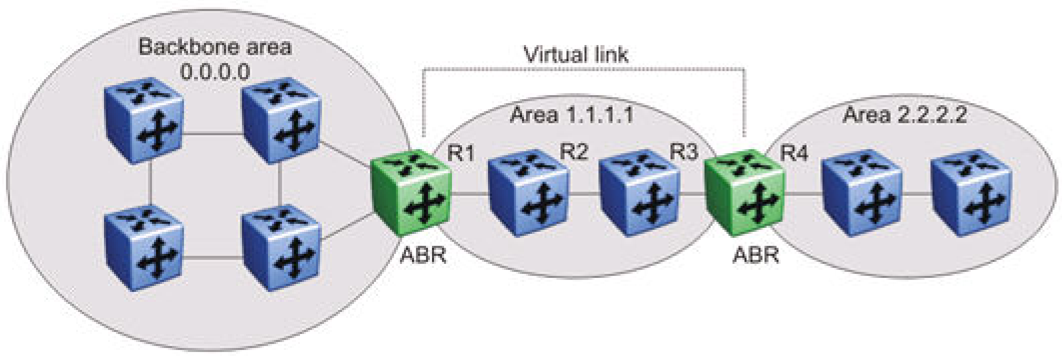

Figure 14: Virtual link between ABRs through a transit area

You can configure automatic or manual virtual links.

An automatic virtual link can provide redundancy support for critical network connections. Automatic virtual linking creates virtual paths for vital traffic paths in your OSPF network. If a connection fails on the network, for example, when an interface cable providing connection to the backbone (either directly or indirectly) becomes disconnected from the switch, the virtual link is available to maintain connectivity.

Specify automatic virtual linking to ensure that a link is created to another router. When you specify automatic virtual linking, this feature is always ready to create a virtual link.

To configure automatic virtual link creation, enable automatic virtual link on both endpoint ABRs (the default value is disabled). Automatic virtual links are removed when the transit area is deleted, auto virtual link is disabled, or the router is no longer an ABR.

If automatic virtual linking uses more resources than you want to expend, a manual virtual link can be the better solution. Use this approach to conserve resources while maintaining specific control of where virtual links are placed in your OSPF network.

To add a virtual link manually, configure both endpoint ABRs with a neighbor router ID and transit area ID. You can configure up to 16 virtual links.

Important:

Auto-created virtual links use default settings that cannot be modified. You can modify parameters for manually added virtual links.

OSPF host route

An OSPF router with hosts directly attached to its interfaces can use host routes to advertise the attached hosts to its neighbors. You can configure up to 32 host routes.

Host routes are identified by the host IP address. You cannot configure the TOS for a host route as TOS-based routing is not supported. For each host directly connected to the router, configure the cost of the link to the host during host creation. You cannot modify this cost.

When a host is added to, or deleted from, a host route, the router updates the router LSAs and floods them to neighbors in each area where that router has an interface.

OSPF interfaces

You can configure an OSPF interface, or link, on an IP interface. On the Avaya Ethernet Routing Switch 4000 Series, an IP interface can be either a brouter port or a VLAN. The system obtains the state information associated with the interface from the underlying lower level protocols and the routing protocol itself.

Important:

To change the interface type of an enabled OSPF interface, you must first disable it, change the type, and then reenable it.

OSPF network types allow OSPF-neighboring between routers over various types of network infrastructures. You can configure each interface to support various network types.

Avaya Ethernet Routing Switch 4000 Series supports the following OSPF network interface type:

- Broadcast interfaces on page 52
- Passive interfaces on page 53

Broadcast interfaces

Broadcast interfaces automatically discover every OSPF router on the network by sending OSPF Hellos to the multicast group AllSPFRouters (224.0.0.5).

Neighboring is automatic and requires no configuration.

Broadcast interfaces support many attached routers and can address a single physical message to all attached broadcast routers (sent to AllSPFRouters and AllDRouters).

Broadcast interfaces dynamically discover neighboring routers using the OSPF Hello Protocol. Each pair of routers on a broadcast network, such as Ethernet, communicate directly.

Passive interfaces

A passive interface is an interfacing network in OSPF that does not generate LSAs or form adjacencies. Passive interfaces are typically used on an access network.

Using passive interfaces limits the amount of CPU cycles required to perform the OSPF routing algorithm.

Use a passive interface to enable an interface to advertise into an OSPF domain while limiting its adjacencies.

When you change the interface type to passive, the interface is advertised into the OSPF domain as an internal stub network with the following behaviors:

- does not send Hello packets to the OSPF domain
- does not receive Hello packets from the OSPF domain
- does not form adjacencies in the OSPF domain

The interface requires only that it be configured as passive to be advertised as an OSPF internal route. If the interface is not a passive interface, to advertise a network into OSPF and not form OSPF adjacencies, the interface must be configured as nonOSPF, and the local network must be redistributed as an autonomous system external (ASE) LSA.

The network behind a passive interface is treated as a stub network and does not form adjacencies. The network is advertised into the OSPF area as an internal route.

OSPF packets

OSPF runs over IP, which means that an OSPF packet is sent with an IP data packet header. The protocol field in the IP header is 89, which identifies it as an OSPF packet.

All OSPF packets start with a 24-octet header that contains information about the OSPF version, the packet type and length, the ID of the router that transmits the packet, and the ID of the OSPF area from which the packet is sent. An OSPF packet is one of the following types:

- Hello packets are transmitted between neighbors and are never forwarded. The Hello Protocol requires routers to send Hello packets to neighbors at pre-defined Hello intervals. A neighbor router that does not receive a Hello packet declares the other router dead.
- Database description (DD) packets are exchanged when a link is established between neighboring routers which synchronize their link-state databases.

- Link-state request packets describe one or more link-state advertisements that a router requests from its neighbor. Routers send link-state requests if the information received in DD packets from a neighbor is not consistent with its own link-state database.
- Link-state update packets contain one or more link-state advertisements and are sent following a change in network conditions.
- Link-state acknowledgement packets are sent to acknowledge receipt of link-state updates and contain the headers of the received link-state advertisements.

OSPF metrics

For OSPF, the best path to a destination is the path that offers the least-cost metric (least-cost delay). OSPF cost metrics are configurable, so you can specify preferred paths. You can configure metric speed globally or for specific interfaces on your network. In addition, you can control redistribution options between non-OSPF interfaces and OSPF interfaces.

Default metric speeds are assigned for different port types, as shown in the following table.

Table 6: OSPF default metrics

Port type	Default OSPF metric
10 Mb/s	100
100 Mb/s	10
1000 Mb/s	1
10 000 Mb/s	1

OSPF security mechanisms

The Avaya Ethernet Routing Switch 4000 implementation of OSPF includes security mechanisms to prevent unauthorized routers from attacking the OSPF routing domain. These security mechanisms prevent a malicious person from joining an OSPF domain and advertising false information in the OSPF LSAs. Likewise, security prevents a misconfigured router from joining an OSPF domain. Currently there are two security mechanisms supported: simple password security and Message Digest 5 (MD5) security.

Simple Password

The Simple Password security mechanism is a simple-text password that is transmitted in the OSPF headers. Only routers that contain the same authentication ID in their LSA headers can communicate with each other.

Important:

Avaya recommends you not to use this security mechanism because the password is stored in plain text, and can be read from the configuration file or from the LSA packet.

Message Digest 5

Avaya recommends that you use Message Digest 5 (MD5) for OSPF security because it provides standards-based (RFC 1321) authentication using 128-bit encryption. When you use MD5 for OSPF security, it is very difficult for a malicious user to compute or extrapolate the decrypting codes from the OSPF packets.

When you use MD5, each OSPF packet has a message digest appended to it. The digest must be matched between sending and receiving routers. The message digest is calculated at both the sending and receiving routers based on the MD5 key and any padding, and then compared. If the message digest computed at the sender and receiver does not match, the packet is rejected.

Each OSPF interface supports up to 2 keys, identifiable by key ID, to facilitate a smooth key transition during the rollover process. Only the selected primary key is used to encrypt the OSPF transmit packets.

Routing Information Protocol

Routing Information Protocol (RIP) is a standards-based, dynamic routing protocol based on the Bellman-Ford (or distance vector) algorithm. It is used as an Interior Gateway Protocol (IGP). RIP allows routers to exchange information to compute the shortest routes through an IPv4-based network. The hop count is used as a metric to determine the best path to a remote network or host. The hop count cannot exceed 15 hops (the distance from one router to the next is one hop).

RIP is defined in RFC 1058 for RIP version 1 and RFC 2453 for RIP version 2. The most significant difference between the two versions is that, while RIP version 1 is classful, RIP version 2 is a classless routing protocol that supports variable length subnet masking (VLSM) by including subnet masks and next hop information in the RIP packet.

RIP Operation

Each RIP router maintains a routing table, which lists the optimal route to every destination in the network. Each router advertises its routing information by sending routing information updates at regular intervals. Neighboring routers use this information to recalculate their routing tables and retransmit the routing information. For RIP version 1, no mask information is exchanged; the natural mask is always applied by the router receiving the update. For RIP version 2, mask information is always included.

RIP uses User Datagram Protocol (UDP) data packets to exchange routing information.

The sequence of processes governed by RIP is as follows:

1. When a router starts, it initializes the RIP data structures and then waits for indications from lower-level protocols that its interfaces are functional.
2. RIP advertisements are sent on all the interfaces that are configured to send routing information.
3. The neighbors send their routing tables and the new router updates its routing table based on the advertisements received.
4. From then on, each router in the network sends periodic updates to ensure a correct routing database.

RIP metrics

RIP is known as a distance vector protocol. The vector is the network number and next hop, and the distance is the cost associated with the network number. RIP identifies network reachability based on cost, and cost is defined as hop count. The distance from one router to the next is considered to be one hop. This cost or hop count is known as the metric.

The following figure shows the hop counts between various units in a network.

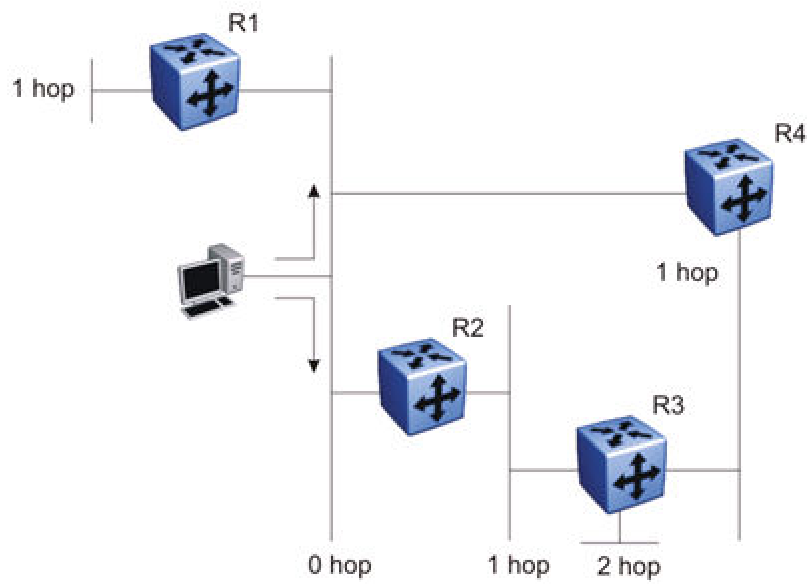

Figure 15: RIP hop counts

A directly connected network has a metric of zero. An unreachable network has a metric of 16. Therefore, 15 hops or 15 routers is the highest possible metric between any two networks.

RIP routing updates

Each RIP router advertises routing information updates out of all RIP-enabled interfaces at regular intervals (30 seconds by default). You can configure this interval using the update timer parameter. The routing updates contain information about known networks and the distances (hop count) associated with each. For RIP version 1, no mask information is exchanged; the natural mask is always applied by the router receiving the update. With RIP version 2, mask information is always included.

If a RIP router does not receive an update from another RIP router within a timeout period (180 seconds by default), it deletes the routes advertised by the nonupdating router from its routing table. You can configure this interval using the timeout interval parameter.

The router keeps aged routes from nonupdating routers temporarily in a garbage list and continues to advertise them with a metric of infinity (16) for a holddown period (120 seconds by default), so that neighbors know that the routes are unreachable. You can configure this interval using the holddown timer parameter. If a valid update for a garbage route is received within the holddown period, the router adds the route back into its routing table. If no update is received, the router completely deletes all garbage list entries for the nonupdating router.

RIP configuration

When the system is switched on, it retrieves the global settings and settings for each interface from the configuration file.

The following global settings are stored in the configuration file:

- Import Metric
- Rip Timer
- Rip State
- Rip Domain
- Timeout
- Holddown

The following interface settings are stored in the configuration file:

- Vlan Id
- Enable
- Advertise When Down
- Auto Aggregation
- Auto Summary
- HoldDown
- In Policy

- Listen
- Out Policy
- Poison
- Proxy Announce
- Rip2 Transmit Mode
- Rip2 Receive Mode
- Triggered Enable
- Rip Out Filter

RIP Features

RIP supports the following standard behavior:

- periodic RIP updates about effective best routes
- garbage collection
- triggered update for changed RIP routes
- broadcast/multicast of regular and triggered updates
- subnet mask (RIP version 2)
- routing table update based on the received RIP message
- global update timer
- holddown timer and timeout timer for each device and interface

The Avaya Ethernet Routing Switch 4000 Series implementation of RIP also supports the following features:

- in and out routing policies
- auto-aggregation (also known as auto-summarization) of groups of adjacent routes into single entries

Many RIP features are configurable. The actual behavior of the protocol depends on the feature configurations.

Virtual Router Redundancy Protocol

The Virtual Router Redundancy Protocol (VRRP) (RFC 3768) can eliminate the single point of failure that can occur when the single static default gateway router for an end station is lost. VRRP allows the use of a virtual IP address (transparent to users) shared between two or more routers connecting a common subnet to the enterprise network. With end hosts using the virtual

IP address as the default gateway, VRRP provides dynamic default gateway redundancy in the event of failure.

VRRP uses the following terms:

- VRRP router: a router running the VRRP protocol.
- Virtual router: the abstract object managed by VRRP that is assigned the virtual IP address and that acts as the default router for a set of IP addresses across a common network. Each virtual router is assigned a virtual router ID (VRID).
- Virtual router master: the VRRP router that assumes responsibility for forwarding packets sent to the IP address associated with the virtual router. The master router also responds to packets sent to the virtual router IP address and answers ARP requests for this IP address.
- Virtual router backup: the router or routers that can serve as the failover router if the master router becomes unavailable. If the master router fails, a priority election process provides a dynamic transition of forwarding responsibility to a new master router.
- Priority: an 8-bit value assigned to all VRRP routers. A higher value represents a higher priority for election to the master router. The priority can be a value from 1 to 255. If two or more switches have the same priority value, the switch with the highest numerical IP address value is selected and becomes the VRRP master. When a master router fails, an election process takes place among the backup routers to dynamically reassign the role of the master router. The host is unaware of the entire process.

VRRP operation

Once you initialize a VRRP router, if there are no other VRRP routers enabled in the VLAN, the initialized router assumes the role of the master router. When additional VRRP routers are enabled in the VLAN, an election process takes place among them to elect a master router, based on their priority.

The master router functions as the forwarding router for the IP address associated with the virtual router. When a host sends traffic to a remote subnet, it sends an ARP request for the MAC address of the default gateway. In this case, the master router replies with the virtual MAC address. The benefit of using a virtual MAC address is that, if the master router fails, the VRRP backup router uses the same virtual MAC address.

The master router responds to ARP requests for the IP address, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts only packets addressed to the IP address associated with the virtual router. The master router also sends VRRP advertisements periodically (every 1 second by default) to all VRRP backup routers.

In the backup state, a VRRP router monitors the availability and state of the master router. It does not respond to ARP requests and must discard packets with a MAC address equal to the virtual router MAC address. It does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, it transitions back to the initialize state.

If the master router fails, the backup router with the highest priority assumes the role of the master router. It transitions to the master state and sends the VRRP advertisement and ARP request as described in the preceding paragraphs. The virtual router IP address and MAC address does not change, providing transparent redundancy.

VRRP topology example

The following figure shows a VRRP topology example.

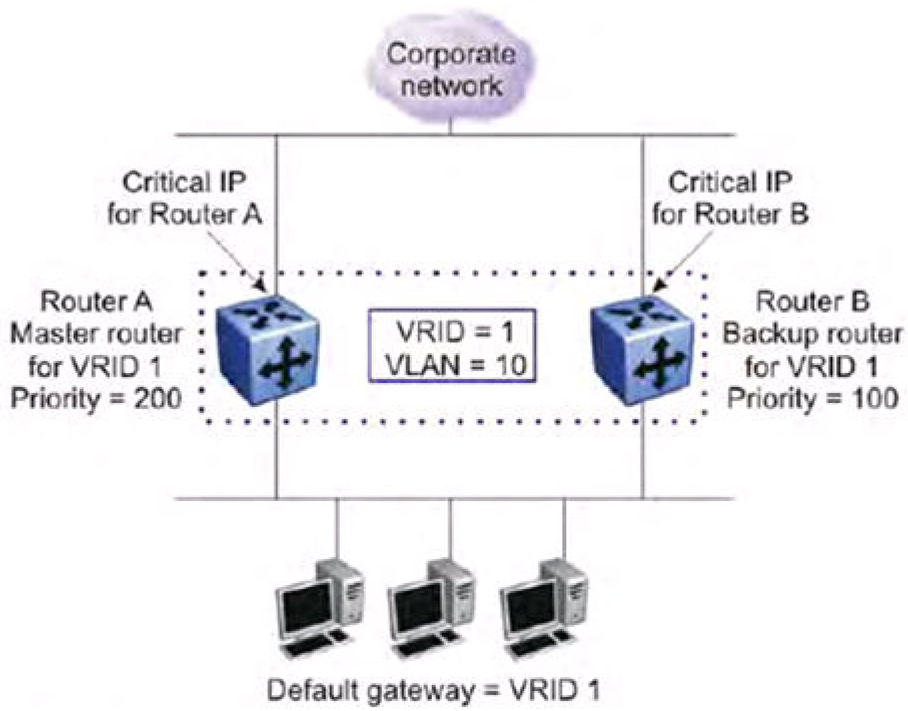

Figure 16: VRRP topology example

In this VRRP example, to configure router A as the master router and router B as the backup router, configure the routers as follows:

1. On router A, create a VLAN, (in this case VLAN 10).
2. Assign an IP address to the VLAN for routing.
3. Configure VRRP properties for VLAN 10 on router A:
    - Assign a virtual router ID (in this case, VRID 1).
    - Set the virtual router IP address to a previously unassigned IP address.
    - Set the priority to a value above the priority of the Router B (in this case, 200).
4. On router B, create a matching VLAN (in this case, VLAN 10).
5. Assign an IP address to the VLAN for routing.
6. Configure VRRP properties for VLAN 10 on router B:
    - Assign the same virtual router ID as on router A (VRID 1).
    - Configure the same virtual router IP address as on router A.
    - Set the priority to a value below that on Router A (in this case, 100).

Once you enable VRRP on both of these switches, an election process takes place, and because router A has the higher priority, it is elected as the master router. It then assumes responsibility for the configured virtual router IP address.

VRRP critical IP address

Within a VRRP VLAN, it is possible for one link to go down, while the remaining links in the VLAN remain operational. Because the VRRP VLAN continues to function, a virtual router associated with that VLAN does not register a master router failure.

As a result, if the local router IP interface connecting the virtual router to the external network fails, this does not automatically trigger a master router failover. The critical IP address resolves this issue. If the critical IP address fails, it triggers a failover of the master router.

You can specify the local router IP interface uplink from the VRRP router to the network as the critical IP address. This ensures that if the local uplink interface fails, VRRP initiates a master router failover to one of the backup routers.

In the **VRRP topology example** figure, the local network uplink interface on router A is shown as the critical IP address for router A. As well, the similar network uplink is shown as the critical IP address for router B. Router B also requires a critical IP address for cases when it assumes the role of the master router.

VRRP fast advertisement interval

With VRRP, you can set the advertisement interval between sending advertisement messages in seconds. This permits faster network convergence with standardized VRRP failover. However, losing connections to servers for more than a second can result in missing critical failures. Customer network uptime in many cases requires faster network convergence, which means network problems must be detected within hundreds of milliseconds.

To meet these requirements the Ethernet Routing Switch supports a fast advertisement interval parameter. The fast advertisement interval is similar to the advertisement interval except for the unit of measure and range. The fast advertisement interval is expressed in milliseconds and the range is from 200 to 1000 milliseconds. To use the fast advertisement interval, you must configure a value for the parameter and explicitly enable the feature.

When the fast advertisement interval is enabled, VRRP can only communicate with other Ethernet Routing Switch devices with the same settings.

Equal Cost Multi Path (ECMP)

With the Equal Cost Multi Path (ECMP) feature, routers can use up to four equal cost paths to the same destination prefix. The L3 switch can use multiple paths for traffic load sharing and in the event of network failure, achieve faster convergence to other active paths. When the L3 switch maximizes load sharing among equal cost paths, the system uses links more efficiently for IP traffic transmission.

Note:

With multiple equal cost paths to a configured network , a route is considered the group of paths (one up to four) to that network, instead of each individual path. This affects “show ip route summary” and “show ip num-routes” outputs that now display the number of groups of equal cost paths to a destination network as the total number of routes.

The ECMP feature supports the following protocols:

- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Static Routes

Route Policies

Using standard routing schemes, a router forwards packets on routes that it has learned through routing protocols such as RIP and OSPF or through the introduction of static routes. With route policies, the router can forward packets based on rule sets created by the network administrator. These rule sets, or policies, are then applied to the learned or static routes.

On the Avaya Ethernet Routing Switch 4000 Series, you can configure route policies for RIP and OSPF. You can use the route policies to perform the following tasks:

- Listen for routing updates from specific gateways.
- Listen for routing updates from specific networks.
- Assign a specific subnet mask to be included with a network in the routing table.
- Advertise routing updates from specific gateways.
- Advertise routing updates to specific networks.
- Assign a specific subnet mask to be included in the route summary packets.
- Advertise routes learned by one protocol to another.

The Avaya Ethernet Routing Switch 4000 Series supports the following types of policies:

- Accept (In) Policies

  Accept polices are applied to incoming routing updates before they are applied to the routing table. In the case of RIP, accept policies can be applied to all incoming packets. Only one policy can be created for each RIP interface. In the case of OSPF, accept policies are only applied to Type 5 External routes based on the advertising router ID. There can only be one OSPF accept policy per switch and the policy is applied before updates are added to the routing table from the link state database.
- Announce (Out) Policies

  Announce policies are applied to outgoing routing updates before the routing update packets are actually transmitted from the switch. In the case of RIP, announce policies can be applied to all outgoing packets. Only one policy can be created for each RIP interface. Announce policies are not supported for OSPF as OSPF requires routing information to be consistent throughout the OSPF domain.

- Redistribution Policies

  Redistribution policies are used to provide notification of addition or deletion of a route in the routing table by one protocol to another protocol. OSPF redistribution policies send redistributed routes as Type 5 External routes. To configure redistribution on a router, it must be an ASBR. There can be only one OSPF redistribution route per switch and redistribution must be enabled. The OSPF accept policy takes precedence over the redistribution policy. You cannot configure a redistribution policy for RIP.

Route policies consist of the following items:

- Prefix-lists
  - List of IP addresses with subnet masks used to define an action
  - Identified by a unique prefix-list name
  - Prefixes, identified by a prefix name, can be created and added in the prefix list using ACLI commands
- Policies
  - Identified by a unique policy name or ID
  - Contains several sequence numbers that in turn contains several significant fields
  - Based on the context of policy usage, the fields are read or ignored; a whole complete policy can be applied to execute a purpose
  - Sequence number also acts as a preference; a lower sequence number has a higher priority
- Routing Protocols
  - Routing Protocol (RP), OSPF and RIP, needs to be registered with the Routing Protocol Server (RPS), and enabled to apply polices. A registered, but disabled RP cannot apply policies. By default, RIP and OSPF, are registered with RPS, and disabled to apply policies
  - RP explicitly informs RPS to send a notification when a specific routing policy object changes; RPS sends a notification message to RP if the requested route policy objects change
  - RP decides whether to re-apply the Accept/Announce policy

Route policies in a stack

In a stacked environment, the following rules apply to routing policies:

- The policy database is stored in all stack units.
- Policy configuration is supported from only the base unit. The base unit sends updates to non-base units to update the policy database in each stack unit.

- During database updates, only the database in the base unit is synchronized with the non-base unit. The database in the non-base units are deleted during the exchange.
- Only the policies stored in the base unit are used by RIP and OSPF for policy application.

Chapter 4: IGMP fundamentals

This chapter provides an overview of IP multicast and Internet Group Management Protocol (IGMP). To support multicast traffic, the Avaya Ethernet Routing Switch 4000 Series provides support for IGMP snooping.

Overview of IP multicast

Most traditional network applications such as Web browsers and e-mail employ unicast connections in which each client sets up a separate connection to a server to access specific data. However, with certain applications such as audio and video streaming, more than one client accesses the same data at the same time. With these applications, if the server sends the same data to each individual client using unicast connections, the multiple connections waste both server and network capacity. For example, if a server offers a 1 Mbit/sec live video stream for each client, a 100 Mbit/sec network interface card (NIC) on the server could be completely saturated after 90 client connections. The following figure shows an example of this waste of resources.

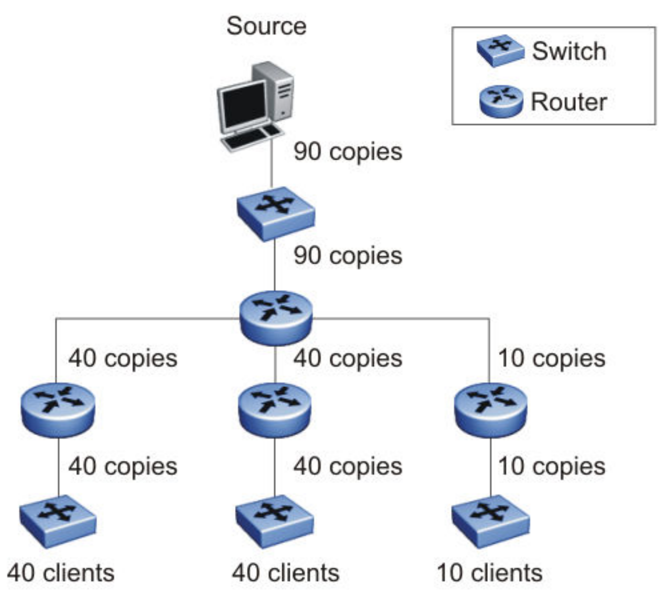

Figure 17: Wasteful propagation of multiple copies of the same unicast stream

Multicasting provides the ability to transmit only one stream of data to all the interested clients at the same time. The following figure shows a simple example of how multicasting works. The source of the multicast data forwards only one stream to the nearest downstream router, and each subsequent downstream router forwards a copy of the same data stream to the recipients who are registered to receive it.

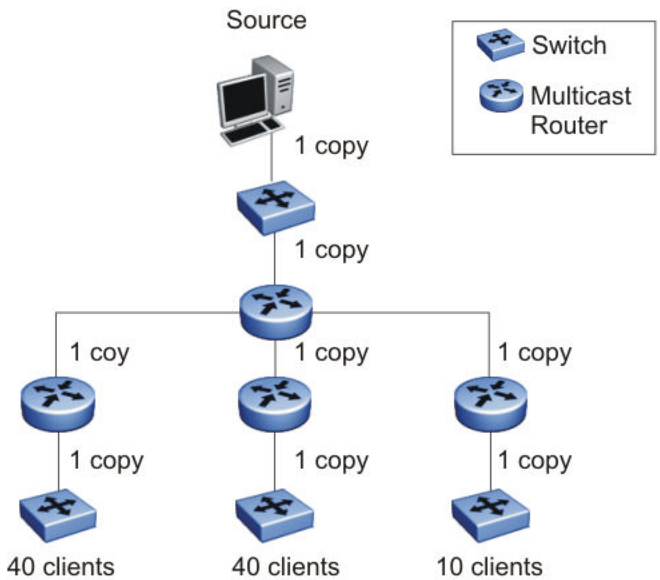

Figure 18: One stream replicated using multicasting

This one-to-many delivery mechanism is similar to broadcasting except that, while broadcasting transmits to all hosts in a network, multicasting transmits only to registered host groups. Because multicast applications transmit only one stream of data, which is then replicated to many receivers, multicasting saves a considerable amount of bandwidth.

Clients that want to receive the stream must register with the nearest multicast router to become a part of the receiving multicast group.

One downside to multicasting is that the multicast streams transmit data using User Datagram Protocol (UDP) packets, which are not as reliable as Transmission Control Protocol (TCP) packets.

Applications that use multicasting to transmit data include the following:

- multimedia conferencing
- real-time data multicasts (such as stock tickers)
- gaming and simulations

Multicast groups

To receive a multicast stream from a particular source, hosts must register with the nearest multicast router. The router adds all interested hosts to a multicast group, which is identified by a multicast IP address.

Multicast routers use Internet Group Membership Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. To identify the hosts that want to be added to a group, a querier router sends out IGMP queries to each local network. A host that wants to belong to the group sends a response in the form of an IGMP membership report.

Each multicast router maintains a multicast routing table that lists each source, group (S,G) pair, which identifies the IP address of the source and the multicast address of the receiving group. For each (S,G) pair, the router maintains a list of downstream forwarding ports to which the multicast traffic is forwarded, and the upstream port where the multicast traffic is received.

Multicast addresses

Each multicast host group is assigned 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.

On the Avaya Ethernet Routing Switch 4000, you cannot use 24-bit subnets like 224.0.0.0/24 and 224.128.0.0/24 for multicast data traffic. This restriction applies to the entire multicast address range from 224.0.0.0/8 to 239.128.0.0/8.

IGMP overview

IGMP is the Layer 3 protocol used by IP multicast routers to learn the existence of multicast group members on their directly attached subnets (see RFC 2236). With IGMP, hosts can register their desired group memberships to their local querier router.

You can configure up to 512 IGMP groups.

A multicast querier router communicates with hosts on a local network by sending IGMP queries. The router periodically sends a general query message to each local network of the router. A host that wants to join a multicast group sends a response in the form of a membership report requesting registration with a group. After the querier router registers hosts to a group, it forwards all incoming multicast group packets to the registered host networks. As long as any host on a subnet continues to participate in the group, all hosts, including nonparticipating end stations on that subnet, receive the IP Multicast stream.

IGMP versions are backward compatible and can all exist together on a multicast network.

The following sections provide more details on the differences between the different IGMP versions.

IGMPv1 operation

IGMP version 1 is the simplest of the IGMP versions and is widely deployed.

IGMPv1 supports the following two message types:

- 0x11 – Membership Query message. Packets are sent to the all-systems multicast group (224.0.0.1).
- 0x12 – Membership Report message. Packets are sent to the group that the host intends to join.

The IGMPv1 router periodically sends host membership queries (also known as general queries) to its attached local subnets to inquire if any hosts are interested in joining any multicast groups. The interval between queries is a configurable value on the router. A host that wants to join a multicast group sends a membership report message to the nearest router, one report for each joined multicast group. After receiving the report, the router adds the Multicast IP address and the host port to its forwarding table. The router then forwards any multicast traffic for that multicast IP address to all member ports.

The router keeps a list of multicast group memberships for each attached network, and a Group Membership Interval timer for each membership. Repeated IGMP membership reports refresh the timer. If no reports are received before the timer expires, the router sends a query message.

In some cases, the host does not wait for a query before it sends report messages to the router. Upon initialization, the host can immediately issue a report for each of the multicast groups that it supports. The router accepts and processes these asynchronous reports the same way it accepts requested reports.

IGMPv1 leave process

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 to determine whether they want to continue to participate. As long as any host on the subnet continues to participate, all hosts, including nonparticipating end stations on the subnet, receive the IP Multicast stream.

If all hosts on the subnet leave the group, the router continues to send general queries to the subnet. If no hosts send reports after three consecutive queries, the router determines that no group members are present on the subnet.

IGMPv2 operation

IGMPv2 extends the IGMPv1 features by implementing a host leave message to quickly report group membership termination to the routing protocol. Instead of routers sending multiple queries before determining that hosts have left a group, the hosts can send a leave message. This feature is important for multicast groups with highly volatile group membership.

The IGMPv2 join process is similar to the IGMPv1 join process.

IGMPv2 also implements a querier election process.

IGMPv2 adds support for the following three new message types:

- 0x11 – General Query and Group Specific Query message.
- 0x16 – Version 2 Membership Report (sent to the destination IP address of the group being reported)
- 0x17 – Version 2 Membership Leave message (sent to all-router [224.0.0.2] multicast address)

IGMPv2 also supports IGMPv1 messages.

Host leave process

With IGMPv2, if the host that issued the most recent report leaves a group, the host issues a leave message. The multicast router on the network then issues a group-specific query to determine whether other group members are present on the network. In the group-specific query message, the Group Address field is the group being queried (the Group Address field is 0 for the General Query message). If no host responds to the query, the router determines that no members belonging to that group exist on that interface.

The following figure shows an example of how IGMPv2 works.

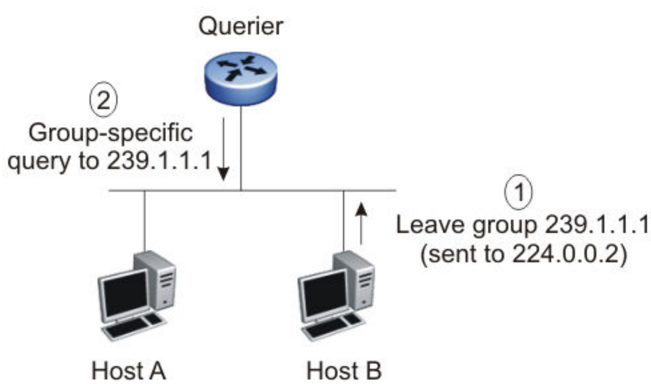

Figure 19: IGMPv2

In this example, the following occurs:

- The host sends a leave message (to 224.0.0.2).
- The router sends a group-specific query to group 239.1.1.1.
- No IGMP report is received.
- Group 239.1.1.1 times out.

Querier election process

Normally only one querier exists per subnet. When multiple IGMPv2 routers are present on a network, the router with the lowest IP address is elected to send queries. All multicast routers start up as a querier on each attached network. If a multicast router receives a query message from a router with a lower IP address, the router with the higher IP address becomes a nonquerier on that network.

IGMP requests for comment

For additional information on IGMP, see the following requests for comment (RFC):

- For IGMPv1, see RFC 1112.
- For IGMPv2, see RFC 2236.
- For IGMP snooping, see RFC 4541.
- For IGMP management information bases (MIB), see RFC 2933.

IGMP snooping

If at least one host on a VLAN specifies that it is a member of a group, by default, the Avaya Ethernet Routing Switch 4000 Series forwards to that VLAN all datagrams bearing the multicast address of that group. All ports on the VLAN receive the traffic for that group.

The following figure shows an example of this scenario. Here, the IGMP source provides an IP Multicast stream to a designated router. Because the local network contains receivers, the designated router forwards the IP Multicast stream to the network. Switches without IGMP snoop enabled flood the IP Multicast traffic to all segments on the local subnet. The receivers requesting the traffic receive the desired stream, but so do all other hosts on the network. Although the nonparticipating end stations can filter the IP Multicast traffic, the IP Multicast traffic still exists on the subnet and consumes bandwidth.

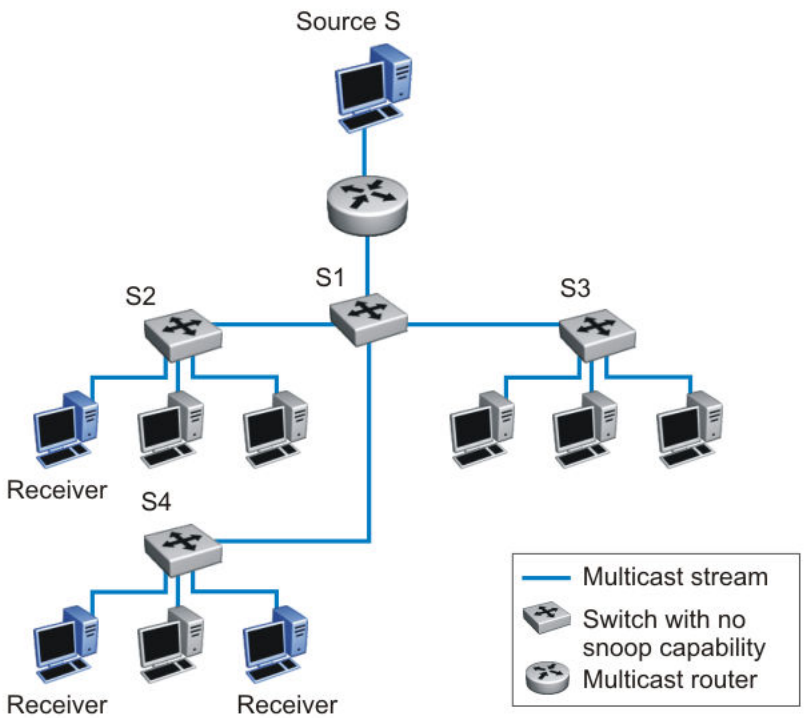

Figure 20: IP multicast propagation on a LAN without IGMP snooping

To prune ports that are not group members from receiving the group data, the Avaya Ethernet Routing Switch 4000 Series supports IGMP snoop for IGMPv1 and IGMPv2. With IGMP snoop enabled on a VLAN, the switch forwards the multicast group data to only those ports that are members of the group. When using IGMP snoop, VLANs can provide the same benefit as IP Multicast routers, but in the local area.

The Avaya Ethernet Routing Switch 4000 identifies multicast group members by listening to IGMP packets (IGMP reports, leaves, and queries) from each port. The switch suppresses the reports by not forwarding them out to other VLAN ports, forcing the members to continuously send their own reports. The switch uses the information gathered from the reports to build a list of group members. After the group members are identified, the switch blocks the IP Multicast stream from exiting any port that does not connect to a group member, thus conserving bandwidth.

As shown in the following figure, after the switches learn which ports are requesting access to the IP Multicast stream, all other ports not responding to the queries are blocked from receiving the IP Multicast data.

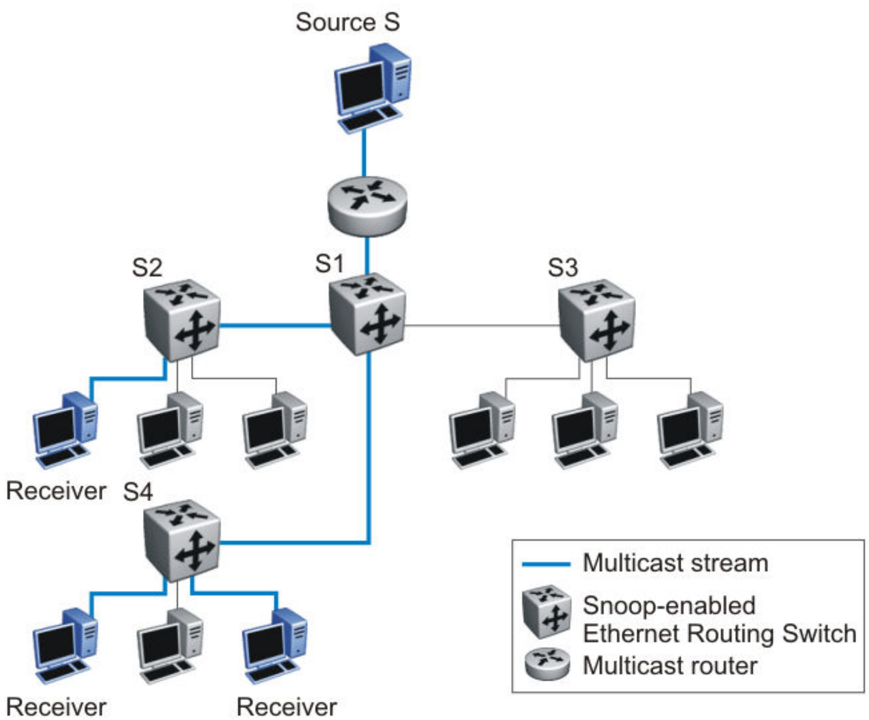

Figure 21: Ethernet Routing Switch running IGMP snooping

The switch continues to forward the IGMP membership reports from the hosts to the multicast routers, and also forwards queries from multicast routers to all port members of the VLAN.

IGMP proxy

With IGMP snoop enabled, the switch can receive multiple reports for the same multicast group. Rather than forward each report upstream, the Ethernet Routing Switch 4000 can consolidate these multiple reports by using the IGMP proxy feature. With IGMP proxy enabled, if the switch receives multiple reports for the same multicast group, it does not transmit each report to the upstream multicast router. Instead, the switch forwards the first report to the querier and suppresses the rest. If new information emerges that another multicast group is added or that a query is received since the last report is transmitted upstream, the report is then forwarded to the multicast router ports.

To enable IGMP Proxy, you must first activate IGMP snooping.

In Figure 22: Ethernet Routing Switch running IGMP proxy on page 73, switches S1 to S4 represent a local area network (LAN) connected to an IP Multicast router. The router periodically sends Host Membership Queries to the LAN and listens for a response from end stations. All of the clients connected to switches S1 to S4 are aware of the queries from the router.

One client, connected to S2, responds with a host membership report. Switch S2 intercepts the report from that port, and generates a proxy report to its upstream neighbor, S1. Also, two clients connected to S4 respond with host membership reports, causing S4 to intercept the reports and to generate a consolidated proxy report to its upstream neighbor, S1.

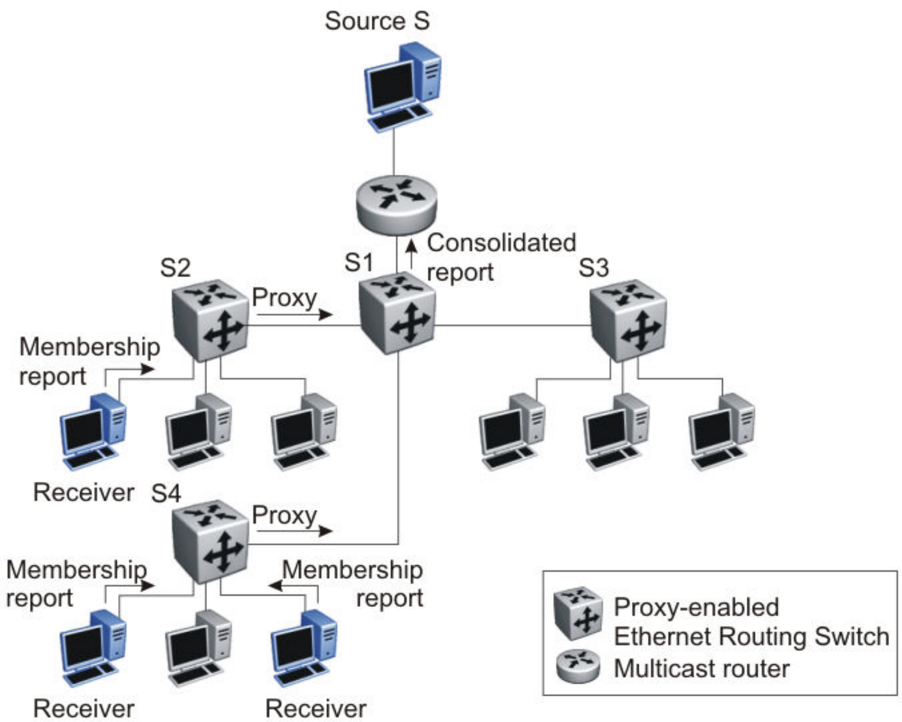

Figure 22: Ethernet Routing Switch running IGMP proxy

Switch S1 treats the consolidated proxy reports from S2 and S4 as if they were reports from any client connected to its ports, and generates a consolidated proxy report to the designated router. In this scenario, the router receives a single consolidated report from that entire subnet.

The consolidated proxy report generated by the switch remains transparent to Layer 3 of the International Standardization Organization, Open Systems Interconnection (ISO/OSI) model. (The switch IP address and Media Access Control [MAC] address are not part of proxy report generation.) The last reporting IGMP group member in each VLAN represents all of the hosts in that VLAN and IGMP group.

Forwarding of reports

When forwarding IGMP membership reports from group members, the Avaya Ethernet Routing Switch 4000 forwards the reports only to those ports where multicast routers are attached. To do this, the switch maintains a list of multicast querier routers and the multicast router (mrouter) ports on which they are attached. The switch learns of the multicast querier routers by listening to the queries sent by the routers where source address is not 0.0.0.0.

Static mrouter port and nonquerier

If two IGMP routers are active on a VLAN, the router with the lower IP address is the querier, and the router with the higher IP address operates as a nonquerier. Only querier routers forward IGMP queries on the VLAN; nonqueriers do not forward IGMP queries. IGMP snoop considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port. IGMP snoop is not aware of nonquerier IGMP routers.

By default, IGMP snoop forwards reports to the IGMP querier router only. To allow the switch to forward reports to the nonquerier router as well, you can configure the port connected to the nonquerier as a static mrouter port.

Figure 23: Static mrouter port and nonquerier on page 74 shows how static mrouter ports operate. In this case, the Avaya Ethernet Routing Switch 4000 has port members 5/1 and 6/1 connected to IGMP routers in VLAN 10. Router 1 is the IGMP querier because it has a lower IP address than router 2. Router 2 is then considered the nonquerier.

By default, the switch learns of the multicast querier routers by listening to the IGMP queries. In this case, port 6/1 connected to querier router 1 is identified as an mrouter port.

To forward reports to IGMP router 2 as well, you can configure port 5/1 on the switch as a static mrouter port. In this case, the IGMP reports are forwarded to both routers.

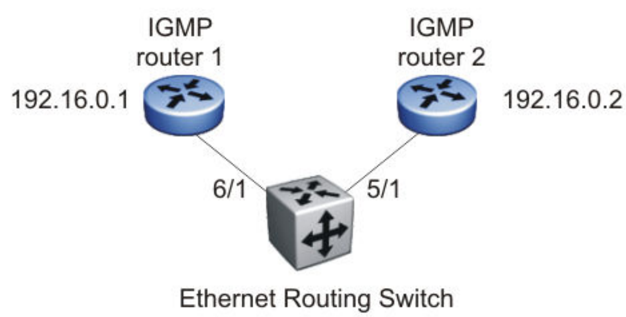

Figure 23: Static mrouter port and nonquerier

Unknown multicast filtering

With IGMP snoop enabled, if the switch receives multicast packets with destination addresses that it has not already registered using IGMP reports, the switch floods all such packets to all ports on the VLAN. All unknown multicast streams of a group are flooded on the VLAN until at least one port in the VLAN becomes a member of that group.

On the Avaya Ethernet Routing Switch 4000, you can enable the unknown multicast filtering feature so that the unknown multicast packets are not flooded on the VLAN. To enable unknown multicast filtering, you can use the **`vlan igmp unknown-mcast-no-flood`** ACLI command.

With this feature enabled, all unknown multicast traffic is forwarded to IGMP mrouter ports only (static and dynamic mrouter ports). If you require unknown multicast traffic to be forwarded to certain ports (for example, to forward Layer 3 multicast routing traffic), set the ports as static mrouter ports.

Avaya recommends that you enable this feature when IGMP snooping is enabled. User settings for the unknown multicast filtering feature are stored in non-volatile random access memory (NVRAM).

Allowing a multicast MAC address to flood all VLANs

The unknown multicast filtering feature introduces a potential problem when a Layer 2 VLAN is placed between two Layer 3 switches that are exchanging protocol packets such as Open Shortest Path First (OSPF). Since the protocols do not join a multicast group, the associated MAC addresses cannot be identified by the IGMP snooping process. These packets are dropped by the Layer 2 switch because the unknown multicast filtering feature is enabled. The two Layer 3 switches can never establish adjacencies, and the OSPF protocol fails.

With the **`vlan igmp unknown-mcast-allow-flood`** ACLI command, you can add MAC addresses that need to be flooded on the switch even when the unknown multicast filtering feature is enabled. The specified MAC addresses are added for all VLANs. Any matching packets are flooded on all ports of a VLAN.

Robustness value

As part of the IGMP snooping configuration, use the robustness value to configure the switch to offset expected packet loss on a subnet. If you expect a network to lose query packets, increase the robustness value.

This value is equal to the number of expected query packet losses for each query interval, plus 1. The range is from 2 to 255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.

IGMP snooping configuration rules

The IGMP snooping feature operates according to specific configuration rules. When configuring your switch for IGMP snooping, consider the following rules that determine how the configuration reacts in any network topology:

- The switch supports up to 512 multicast groups.

  If the multicast group table reaches its limit, a new entry cannot be added with a JOIN message or a new sender identifying a new group. The multicast stream from the new sender is discarded by the hardware. New entries can be added again when the table is not full.

- You cannot configure port mirroring on a static mrouter port.
- If you configure a Multi-Link Trunk member as a static mrouter port, all the Multi-Link Trunk members become static mrouter ports. Also, if you remove a static mrouter port that is a Multi-Link Trunk member, all Multi-Link Trunk members are automatically removed as static mrouter port members.
- Ports must belong to the VLAN on which they are configured as static mrouter ports.
- When Spanning Tree is enabled, the switch learns IGMP groups only on ports that are *not* in Listening or Blocking Spanning Tree states (or, when in RSTP/MSTP mode, only on ports that are in the Designated state). The switch also learns the groups if STP is disabled on a port.
- The IGMP snooping feature is not Rate Limiting-dependent.
- Enabling igmp proxy without having enabled igmp snooping will enable both snooping and proxy. However trying to disable snooping with proxy enabled will produce an error message.
- During any transition from standalone mode to stack mode (or vice versa), the switch deletes all IGMP interfaces that were previously learned and active.

Important:

Because IGMP snooping is set up per VLAN, all IGMP changes are implemented according to the VLAN configuration for the specified ports. It is no longer necessary to specify an mrouter per igmp version, the new syntax permits the configuration of an mrouter port from VLAN configuration port without the need to specify the mrouter port version(the option is unavailable in the new syntax).

Default IGMP values

The following table lists the default IGMP values:

Table 7: Default IGMP values

Parameters	Range	Default Value
Snooping	Enable/Disable	Disable
Version	1-3	2
Proxy	Enable/Disable	Disable
Query Interval	0-65535	125
Robustness Value	2-255	2

IGMP snooping interworking with Windows clients

This section describes an interworking issue between Windows clients and the Ethernet Routing Switches when IGMP snoop is enabled for multicast traffic.

Under normal IGMP snoop operation, as soon as a client joins a specific multicast group, the group is no longer unknown to the switch, and the switch sends the multicast stream only to the ports which request it.

Windows clients, in response to IGMPv2 queries from the switch, reply with IGMPv2 reports.

To force a Windows client to only use IGMPv1 or IGMPv2 reports, change the TCP/IP settings in the Windows Registry located under the following registry key:

```
HKEY_LOCAL_MACHINE
\SYSTEM
\CurrentControlSet
\Services
\Tcpip
\Parameters
```

The specific parameter which controls the IGMP Version is:

```
IGMPVersion
Key: Tcpip\Parameters
Value Type: REG_DWORD—Number
Valid Range: 2, 3, 4
Default: 4
```

To set the Windows Client to only utilize IGMPv2, change the IGMPVersion parameter to 3 (2 specifies IGMPv1, 3 specifies IGMPv2, and 4 specifies IGMPv3).

The IGMPVersion parameter may not be present in the list of the TCP/IP parameters. By default, the system assumes the IGMPv3 value (4). To configure the system for IGMPv2, create the parameter as a DWORD key in the registry and specify Decimal 3.

Important:

If you edit the Windows registry incorrectly, you can severely damage your system. As a minimal safeguard, back up your system data before undertaking changes to the registry.

IGMPv3 snooping

IGMPv3 provides the ability to pack multiple group members in a single Report message, hence reducing the amount of network traffic. Also, IGMPv3 allows a host to include or exclude a list of source addresses for each multicast group of which the host is a member. Routers merge the source address requirements of different hosts for each group.

The Ethernet Routing switch 4000 Series switch supports IGMPv3 source filtering capability with IGMPv3 Snooping. IGMPv3 Snooping remains backward compatible with IGMPv1 and IGMPv2.

IGMP Querier

A multicast query router communicates with hosts on a local network by sending IGMP queries. This router periodically sends a general query message to each local network of the router. This is standard multicast behavior.

It is recommended that each VLAN using IGMP multicast have a router performing multicast queries. This router typically has PIM or DVMRP enabled. Networks with no standalone devices currently have no capability for implementing the pruning of IGMP traffic. The IGMP Querier functionality allows a switch or stack to be configured as an active query router without the need for dedicating a standalone switch in each network to the task.

There are several behavioral differences between a traditional query router and a switch or stack using the IGMP Querier functionality. The following differences should be noted:

- There is no election process. When a switch or stack restarts, the code will send some queries as part of IGMP start up. This process will stop other devices sending queries while they detect the new device starting up. The last active device sending queries on the network is the active one. This is not the case with Layer 3 IGMP behavior.
- If the current active device stops sending queries, a timeout period must elapse before another device takes over. This may result in an ageout of groups, and subsequent flooding, before a new query is sent and the pruning process restarts. This occurs only during the transition between active query devices. Once the new device is established, queries will be sent as configured in the Query Interval and Robust Values fields.
- Multiple active query devices are not supported. Enabling multiple devices establishes one active device and other devices listening to take over should the active device fail.

IGMP Querier functionality can only be enabled when IGMP snooping is active on the switch or stack.

When IGMP snooping send-query is enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switch/host that wants to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

Successful deployment of this feature is dependent on the addition of IP addresses from all devices in the IGMP domain. This is true even when non-management VLANs are used.

IGMP Selective Channel Block

IGMP Selective Channel Block gives you the control to block the streaming of specific channels on some ports.

In certain deployment scenarios, you may prefer to disallow the multicast streaming from specific group addresses to users on specific ports. With IGMP selective channel block feature, you can configure the IGMP membership of ports by blocking IGMP reports received from users on that port, destined for the specific group address/ addresses. The filter can be configured to block a single multicast address or range of addresses.

IGMP Selective Channel Block works regardless of whether the switch is in Layer 2 IGMP snooping mode or the full IGMP mode as the blocking of channels is implemented by blocking the ports from joining an IGMP group. It will also be applicable for IGMP v1, v2 and v3.

You can configure up to 240 channels for blocking.

You cannot use this feature to snoop the multicast streams that are sent from a group to a port.

You can use IGMP Selective Channel Block for both MLT and LACP trunk interfaces. You cannot apply profiles directly to MLT/LACP trunks as you need to apply the profile to a member of the trunk.

When you apply a profile to a port, which belongs to a MLT or LACP trunk, the system applies the profile to all ports of the MLT or LACP. When you dynamically add or remove a port from a MLT or LACP which has a profile associated with it; then the system adds or removes all ports from the profile.

You can use IGMP Selective Channel Block in the standalone as well as in the stacking mode. In stacking mode, the configuration propagates from any unit to all the other units.

Chapter 5: IP routing configuration using ACLI

This chapter describes the procedures you can use to configure routable VLANs using the ACLI.

The Avaya Ethernet Routing Switch 4000 Series are Layer 3 switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address is attached to the VLAN. When routing is enabled in Layer 3 mode, every Layer 3 VLAN is capable of routing and carrying the management traffic. You can use any Layer 3 VLAN instead of the Management VLAN to manage the switch.

It is not a prerequisite to enable global IP routing before configuring an IP address on a VLAN interface. You can configure all IP routing parameters on the switch before you enable routing. When you assign an IP address to the VLAN or brouter port, the system automatically enables routing on the specified VLAN. You must enable global IP routing for the system to route L3 traffic between VLAN interfaces.

For more information about creating and configuring VLANs, see *Configuring VLANs, Spanning Tree, and Multi-Link Trunking on Avaya Ethernet Routing Switch 4000 Series*, NN47205-501.

Configuring global IP routing status

Use this procedure to enable and disable global routing at the switch level. By default, routing is disabled.

Procedure steps

To configure the status of IP routing on the switch, enter the following from the Global Configuration mode:

```
[no] ip routing
```

Variable definitions

The following table describes the **`ip routing`** command variables.

Variable	Value
`no`	Disables IP routing on the switch.

Displaying global IP routing status

Use this command to display the status of IP routing on the switch.

Procedure steps

To display the status of IP routing on the switch, enter the following from the User EXEC mode:

```
show ip routing
```

Dynamic Routing Table allocation configuration using ACLI

The following section provides procedures you can use to manually assign IPv4 route table allocation on the switch.

Configuring Dynamic Routing Total Routes using ACLI

Use the following procedure to manually partition route entries between the routing protocols.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure Dynamic Routing Total Routes by using the following command:

```
ip num-routes {max-local <2-256>} {max-static <0-256>}
```

Variable definitions

The following table defines parameters that you can enter with the **ip num-routes {max-local <2-256>} {max-static <0-256>}** command.

Variable	Value
max-local *<2-256>*	Specifies the maximum number of local routes. Values range from 2 to 256.
max-static *<0-256>*	Specifies the maximum number of static routes. Values range from 0 to 256.

Important:

If you increase the combined maximum totals of local and static routes, the total number of available dynamic routes is decreased. You can decrease the total number of available dynamic routes only when IP routing is disabled. For more information, see Dynamic Routing Table Allocation on page 30.

Configuring Dynamic Routing Total Routes to default using ACLI

Use the following procedure to use default values to manually partition route entries between the routing protocols.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure Dynamic Routing Total Routes to default by using the following command:

```
default ip num-routes [max-local] [max-static]
```

Variable definitions

The following table defines optional parameters that you can enter after the **default ip num-routes** command.

Variable	Value
`max-local`	Configures the maximum number of local routes to the default value (64).
`max-static`	Configures the maximum number of static routes to the default value (32).

Viewing Dynamic Routing Total Routes information using ACLI

Use the following procedure to display the Dynamic Routing Total Routes configuration in the IP Route Resources Partitioning Table.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

View Dynamic Routing Total Routes information by using the following command:

```
show ip num-routes
```

Job aid: show ip num-routes command output

The following figure displays sample output for the **`show ip num-routes`** command.

```
ERS-4524GT>show ip num-routes
===================================================
      IP Route Resources Partitioning Table
===================================================
TYPE                 MAXIMUM  ALLOCATED  AVAILABLE
---------------------------------------------------
Local routes    :         3          1          2
Static routes   :        13          1         12
Dynamic routes  :       496          0        496
---------------------------------------------------
Total routes    :       512          2        510
ERS-4524GT>
```

Figure 24: show ip num-routes command output

Configuring an IP address for a VLAN

To enable routing an a VLAN, you must first configure an IP address on the VLAN.

Procedure steps

To configure an IP address on a VLAN, enter the following from the VLAN Interface Configuration mode:

```
[no] ip address <ipaddr> <mask> [<MAC-offset>]
```

Variable definitions

The following table describes the **ip address** command variables.

Variable	Value
[no]	Removes the configured IP address and disables routing on the VLAN.
<ipaddr>	Specifies the IP address to attach to the VLAN.
<mask>	Specifies the subnet mask to attach to the VLAN
[*<MAC-offset>*]	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address. The valid range is 1-256. Specify the value 1 for the Management VLAN only. If no MAC offset is specified, the switch applies one automatically.

Configuring IP routing status on a VLAN

Use this procedure to enable and disable routing for a particular VLAN.

Procedure steps

To configure the status of IP routing on a VLAN, enter the following from the VLAN Interface Configuration mode:

```
[default] [no] ip routing
```

Variable definitions

The following table describes the **ip routing** command variables.

Variable	Value
`default`	Disables IP routing on the VLAN.
`no`	Disables IP routing on the VLAN.

Displaying the IP address configuration and routing status for a VLAN

Use this procedure to display the IP address configuration and the status of routing on a VLAN.

Prerequisites

Use this command in the Privileged EXEC mode.

Procedure steps

Enter the following command:

`show vlan ip [id <1-4094>]`

Variable definitions

The following table describes the parameters for the **`show vlan ip`** command.

Variable	Value
`id <1-4094>`	Specifies the VLAN ID of the VLAN to be displayed. RANGE: 1-4094.

Job aid

The following table shows the field descriptions for the **`show vlan ip`** command.

Field	Description
Vid	Specifies the VLAN ID.
ifIndex	Specifies an Index entry for the interface.
Address	Specifies the IP address associated with the VLAN.

Field	Description
Mask	Specifies the mask.
MacAddress	Specifies the MAC address associated with the VLAN.
Offset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address.
Routing	Specifies the status of routing on the VLAN: enabled or disabled.

Displaying IP routes

Use this procedure to display all active routes on the switch.

Procedure steps

To display IP routes, enter the following from the User EXEC command mode:

`show ip route [<dest-ip>] [-s <subnet> <mask>] [static] [summary]`

Variable definitions

The following table describes the **show ip route** command variables.

Variable	Value
`<dest-ip>`	Specifies the destination IP address of the routes to display.
`[-s <subnet> <mask>]`	Specifies the destination subnet of the routes to display.
`static`	Displays static route information.
`summary`	Displays a summary of IP route information.

Job aid

The following table shows the field descriptions for the **show ip route** command.

Field	Description
DST	Identifies the route destination.
MASK	Identifies the route mask.
NEXT	Identifies the next hop in the route.
COST	Identifies the route cost.
VLAN	Identifies the VLAN ID on the route.
PORT	Specifies the ports.
PROT	Specifies the routing protocols. For this release, options are LOC (local route) or STAT (static route).
TYPE	Indicates the type of route as described by the Type Legend in the ACLI command display.
PREF	Specifies the route preference.

Brouter port configuration using ACLI

This section provides procedures you can use to configure brouter ports for the ERS 4000 switches.

Configuring a brouter port using ACLI

Use this procedure to create and manage a brouter port on the switch.

Prerequisites

- Use this command in the Ethernet Interface Configuration mode.

Procedure steps

Enter the following command:

```
brouter [port <brouter_port>] vlan <vid> subnet <ip_address/
mask> [routing enable]
```

Variable definitions

The following table describes the parameters for the **brouter** command.

Variable definition

Variable	Value
port <*brouter_port*>	Specifies the port to configure as a brouter port.
vlan <*vid*>	Specifies the VLAN ID of the brouter. When creating a new brouter port, this is the VLAN ID assigned to the brouter port.
subnet <*ip_address/mask*>	Specifies the IP address and subnet mask of the brouter. When creating a new brouter, this is the IP address and subnet mask assigned. RANGE: Subnet mask - 0 to 32
[routing enable]	Enables Layer 3 routing on the brouter port.

Displaying the brouter port configuration using ACLI

Use this procedure to display the brouter port configuration on the switch.

Prerequisites

- Use this command in the User Exec mode.

Procedure steps

To display the brouter port configuration, enter the following from the User Exec mode:

```
show brouter [port <brouter_port>]
```

Variable definitions

The following table describes the parameters for the show brouter command.

Variable definition

Variable	Value
port<*brouter_port*>	Specifies a specific brouter port to be displayed. If you do not use this parameter, the command displays all brouter ports.

Modifying the brouter port IP address using ACLI

Use this procedure to modify the IP address for the brouter port on the switch.

Prerequisites

- Use this command in the Ethernet Interface Configuration mode.

Procedure steps

Enter the following command:

```
brouter [port <brouter_port>] subnet <ip_address/mask>
```

Variable definitions

The following table describes the parameters for the **brouter** command.

Variable definition

Variable	Value
port <brouter_port>	Specifies a specific brouter port to be modified. If you do not use this parameter, the command modifies the brouter port specified in the interface Ethernet <brouter_port> command.
subnet <ip_address/mask>	Specifies the IP address and subnet mask of the brouter. When modifying a brouter port, this is the new IP address and subnet mask to assign to the port. RANGE: subnet mask – 0 to 32

Deleting the brouter port using ACLI

Use this procedure to delete the brouter port on the switch.

Prerequisites

- Use this command in the Ethernet Interface Configuration mode.

Procedure steps

Enter the following command:

```
no brouter [port <brouter_port>]
```

Variable definitions

The following table describes the parameters for the **no brouter** command.

Variable definition

Variable	Value
no	To delete the brouter
port <brouter_port>	Specifies a specific brouter port to be deleted. If you do not use this parameter, the command deletes the brouter port specified in the interface Ethernet <brouter_port> command.

Disabling IP routing for the brouter port using ACLI

Use this procedure to disable IP routing for the brouter port on the switch.

Prerequisites

- Use this command in the Ethernet Interface Configuration mode.

Procedure steps

Enter the following command:

```
no brouter [port <brouter_port>] routing enable
```

Variable definitions

The following table describes the parameters for the **no brouter routing enable** command.

Variable definition

Variable	Value
no	To disable IP routing for the brouter port
port <brouter_port>	Specifies a specific brouter port to be modified. If you do not use this parameter, the command disables IP routing on the brouter port specified in the interface Ethernet <brouter_port> command.
routing enable	Designates Layer 3 routing on the brouter port.

Chapter 6: Static route configuration using ACLI

This chapter describes the procedures you can use to configure static routes using the ACLI.

Configuring a static route

Create static routes to manually configure a path to destination IP address prefixes.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLANs to be routed.

Procedure steps

To configure a static route, enter the following from the Global Configuration command mode:

```
[no] ip route <dest-ip> <mask> <next-hop> [<cost>] [disable]
[enable] [weight <cost>]
```

Variable definitions

The following table describes the **ip route** command variables.

Variable	Value
[no]	Removes the specified static route.
<dest-ip>	Specifies the destination IP address for the route being added. 0.0.0.0 is considered the default route.

Variable	Value
<mask>	Specifies the destination subnet mask for the route being added.
<next-hop>	Specifies the next hop IP address for the route being added.
[*<cost>*]	Specifies the weight, or cost, of the route being added. Range is 1-65535.
[enable]	Enables the specified static route.
[disable]	Disables the specified static route.
[weight *<cost>*]	Changes the weight, or cost, of an existing static route. Range is 1-65535.

Displaying static routes

Use this procedure to display all static routes, whether these routes are active or inactive.

Procedure steps

To display a static route, enter the following from the User EXEC command mode:

`show ip route static [<dest-ip>] [-s <subnet> <mask>]`

Variable definitions

The following table describes the **show ip route static** command variables.

Variable	Value
<dest-ip>	Specifies the destination IP address of the static routes to display.
[-s *<subnet>* *<mask>*]	Specifies the destination subnet of the routes to display.

Job aid

The following table shows the field descriptions for the **show ip route static** command.

Field	Description
DEST	Identifies the route destination.
MASK	Identifies the route mask.
NEXT	Identifies the next hop in the route.
COST	Identifies the route cost.
PREF	Specifies the route preference.
LCLNHOP	Specifies the local next hop status.
STATUS	Specifies the static route status. Options are ACTIVE (in use and present in routing table) or INACTV (not in use and not present in routing table).
ENABLE	Specifies the administrative state of the static route. Options are TRUE (administratively enabled) or FALSE (administratively disabled).

Configuring a management route

Use this procedure to create a management route to the far end network, with a next-hop IP address from the management VLAN's subnet. You can configure a maximum of four management routes on the switch.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the management VLAN interface.

Procedure steps

To configure a static management route, enter the following from the Global Configuration command mode:

```
[no] ip mgmt route <dest-ip> <mask> <next-hop>
```

Variable definitions

The following table describes the **ip mgmt route** command variables.

Variable	Value
[no]	Removes the specified management route.
<dest-ip>	Specifies the destination IP address for the route being added.
<mask>	Specifies the destination subnet mask for the route being added.
<next-hop>	Specifies the next hop IP address for the route being added.

Displaying the management routes

Use this procedure to display the static routes configured for the management VLAN.

Procedure steps

To display the static routes configured for the management VLAN, enter the following from the User EXEC mode:

```
show ip mgmt route
```

Job aid

The following table shows the field descriptions for the **show ip mgmt route** command.

Field	Description
Destination IP	Identifies the route destination.
Subnet Mask	Identifies the route mask.
Gateway IP	Identifies the next hop in the route.

Chapter 7: OSPF configuration using ACLI

This chapter describes the procedures you can use to configure OSPF using ACLI.

The Open Shortest Path First (OSPF) Protocol is an Interior Gateway Protocol (IGP) that distributes routing information between routers belonging to a single autonomous system (AS). Intended for use in large networks, OSPF is a link-state protocol which supports IP subnetting and the tagging of externally-derived routing information.

OSPF commands used during the configuration and management of VLANs in the Interface Configuration mode can be used to configure any VLAN regardless of the one used to log into the command mode. Insert the keyword **`vlan`** with the number of the VLAN to be configured after the command keywords **`ip ospf`**. The current VLAN remains the one used to log into the Interface Configuration command mode after the command execution.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN that you want to enable with OSPF.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

Enabling OSPF globally

Use this procedure to enable OSPF globally on the switch. OSPF is disabled by default.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure OSPF globally on the switch by using the following command:

```
[default] [no] router ospf enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Disables OSPF globally on the switch.
[no]	Disables OSPF globally on the switch.
enable	Enables OSPF globally on the switch. If omitted, enters OSPF Router configuration mode without enabling OSPF.

Configuring the router ID

Use this procedure to configure the router ID, which is expressed in the form of an IP address.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure the router ID by using the following command:

```
[no] router-id <router_id>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Resets the router ID to 0.0.0.0.
<router_id>	Specifies the unique identifier for the router.

Configuring the OSPF default cost metric

Use this procedure to define the OSPF default cost metric.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure the OSPF default cost metric by using the following command:

```
[default] [no] default-cost {ethernet | fast-ethernet | gig-ethernet | ten-gig-ethernet} <metric_value>
```

Variable definitions

The following table describes the command variables.

Variable	Value
`[default]`	Sets the OSPF default cost metric to factory default values. The default values are as follows: • ethernet (10 Mb/s): 100 • fast-ethernet (100 Mb/s): 10 • gig-ethernet (1000 Mb/s): 1 • ten-gig-ethernet (10000 Mb/s): 1
`<metric_value>`	Specifies the default cost metric to assign to the specified port type. The metric value is an integer between 1 and 65535.

Configuring OSPF RFC 1583 compatibility

Use this procedure to OSPF RFC 1583 compatibility.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure OSPF RFC 1583 compatibility by using the following command:

```
[default] [no] rfc1583-compatibility enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
`[default]`	Sets OSPF RFC 1583 compatibility to the default value (enabled).
`[no]`	Disables OSPF RFC 1583 compatibility.

Configuring the OSPF hold down timer

Use this procedure to configure the OSPF hold down timer.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure the OSPF hold own timer by using the following command:

```
[default] timers basic holddown <timer_value>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets the hold own timer to the default value.
<timer_value>	Specifies a hold down timer value between 3 and 60 seconds.

Enabling OSPF system traps

Use this procedure to enable OSPF system traps.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Enable OSPF system traps by using the following command:

```
[no] [default] trap enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets OSPF system traps to the default value (disabled).
[no]	Disables OSPF system traps.

Displaying global OSPF parameters

Use this procedure to display global OSPF parameters.

Procedure steps

To display global OSPF parameters, enter the following from the User EXEC command mode:

```
show ip ospf
```

Configuring OSPF area parameters

Use this procedure to configure OSPF area parameters.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure the OSPF area parameters by using the following command:

```
[default] [no] area <area-id>
[default-cost {0-16777215}]
[import {external | noexternal | nssa}]
[import-summaries {enable}]
[range {ip_addr/subnet_mask} {nssa-entlink | summary-link}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets the specified parameter to the default value (applicable only for default-cost, import, import-summaries, and range).
[no]	Removes the specified OSPF configuration (applicable only for import-summaries [disables] and range [removes the specified range]).
<area-id>	Specifies the Area ID in dotted decimal notation (A.B.C.D).
default-cost *{0-16777215}*	Specifies the default cost associated with an OSPF stub area.
import *{external | noexternal | nssa}*	Specifies the area type by defining the area's support for importing Autonomous System external link state advertisements: • external: specifies a normal area • noexternal: specifies a stub area • nssa: specifies an NSSA **Important:** The configuration of a totally stubby area (no summary advertising) is a two step process. First, define an area with the import flag set to *noexternal*. Second, disable import summaries in the same area with the command **no area <area-id> import-summaries enable**.
import-summaries *{enable}*	Controls the import of summary link state advertisements into stub areas. This setting has no effect on other areas.
range {ip_addr/ subnet_mask} [{nssa-entlink | summary-link}]	Specifies range parameters for the OSPF area.

Displaying OSPF area configuration

Use this procedure to display OSPF area configuration.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF area configuration by using the following command:

```
show ip ospf area [<area-id>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
<area-id>	Displays configuration information about the specified OSPF area. Omitting this parameter displays information for all OSPF areas.

Displaying OSPF area range information

Use this procedure to display OSPF area range information.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF area range information by using the following command:

```
show ip ospf area-range [<range>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
<range>	Displays configuration information about the specified OSPF area range. Omitting this parameter displays information for all OSPF area ranges.

Enabling OSPF on an IP interface

Use this procedure to enable OSPF routing on an IP interface.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Enable OSPF on an interface by using the following command:

```
network <ip_address> [area <area_id>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables OSPF routing on an interface.
<ip_address>	Specifies the IP address of interface to be enabled for OSPF routing.
area <area_id>	Specifies the ID of the area assigned to the interface in dotted decimal notation (A.B.C.D).

Assigning an interface to an OSPF area

Use this procedure to assign an interface to an OSPF area.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Assign an interface to an OSPF area by using the following command:

```
ip ospf area <area-id>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<area-id>	Specifies the unique ID of the area to which the interface connects. An area ID of 0.0.0.0 indicates the OSPF area backbone and is created automatically by the switch.

Configuring OSPF for an interface

Use this command to configure OSPF for an interface.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure OSPF for an interface by using the following command:

```
ip ospf [advertise-when-down enable] [area <A.B.C.D>]
[authentication-key <WORD>] [authentication-type <message-
digest>|<none>|<simple>] [cost <interface_cost>] [dead-
interval <interval>] [enable] [hello-interval <interval>] [mtu-
ignore enable] [network <broadcast | passive>] [port <LINE>]
[primary-md5-key <1-255>] [priority <0-255>] [retransmit-
interval <1-3600>] [transit-delay <1-3600>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
`advertise-when-down enable`	Enables the advertisement of the OSPF interface, and even if the port or VLAN for the routing interface subsequently goes down, the switch continues to advertise the route. **Note:** If a port or VLAN is not operational for the routing interface, no advertisement occurs, even if you enable the *advertise-when-down* parameter.
`authentication-key` *<WORD>*	Specifies an alphanumeric authentication key for the interface. The authentication key can be a maximum of 8 characters.
`authentication-type` *<message-digest>*|*<none>*|*<simple>*	Specifies the type of authentication for the interface. Values include: • *message-digest*: MD5 digest authentication type • *none*: no authentication type is applied to the interface • *simple*: simple password authentication type DEFAULT: none

Variable	Value
cost *<interface_cost>*	Specifies the cost assigned to the interface. This is an integer value between 1 and 65535.
dead-interval *<interval>*	Specifies a dead interval for the interface. This is the interval of time that a neighbor waits for a Hello packet from this interface before the neighbor declares it down. This is an integer value between 0 and 2147483647.
enable	Enables OSPF for the interface. DEFAULT: disabled
hello-interval *<interval>*	Specifies the amount of time between transmission of hello packets from this interface. This is an integer value between 1 and 65535.
mtu-ignore enable	Instructs the interface to ignore the packet MTU size specified in Database Descriptors.
network {broadcast | passive}	Defines the type of OSPF interface this interface is.
port *<LINE>*	Specifies an alternate switch port or list of switch ports for which to configure OSPF. **Note:** This parameter is not available in VLAN Interface Configuration mode.
primary-md5-key *<1-255>*	Specifies the primary MD5 key value to use for authentication. Values range from 1 to 255.
priority *<0-255>*	Assigns a priority to the interface for the purposes of Designated Router election. This is an integer value between 0 and 255.
retransmit-interval *<1-3600>*	Defines the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. This is an integer value between 1 and 3600.
transit-delay *<1-3600>*	Defines the transit delay for this OSPF interface in seconds. The transit delay is the estimated number of seconds it takes to transmit a link-state update over the interface. This is an integer value between 1 and 3600.

Displaying OSPF interface timers

Use this procedure to display OSPF interface timers

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF timers for an interface by using one of the following commands:

```
show ip ospf timer interface [vlan <vid>]
```

OR

```
show ip ospf int-timers
```

Variable definitions

The following table describes the command variables.

Variable	Value
vlan *<vid>*	Displays configured timers for the specified VLAN.

Displaying OSPF timers for virtual links

Use this procedure to display OSPF timers for virtual links.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF timers for virtual links by using the following commands:

```
show ip ospf timer virtual-links
```

Displaying OSPF interface configurations

Use this procedure to display OSPF interface configurations.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF interface configurations by using the following command:

`show ip ospf interface vlan <vid>`

Variable definitions

The following table describes the command variables.

Variable	Value
`vlan <vid>]`	Displays OSPF configuration for the specified interface. If no interface is specified, all interface configurations are displayed.

Displaying OSPF neighbors

Use this procedure to display information about OSPF neighbors for the router.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF neighbors by using the following command:

```
show ip ospf neighbor
```

Specifying a router as an ASBR

Use this procedure to identify a router as an Autonomous System Boundary Router (ASBR).

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure a router as an ASBR by using the following command:

```
[default] [no] as-boundary-router [enable]
```

Variable definitions

The following table describes the command variables.

Variable	Value
`default`	Configures ASBR for the switch to the default value (disabled).
`no`	Disables ASBR for the switch.

Configuring the OSPF authentication type for an interface

Use this procedure to configure the interface authentication type.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure the interface authentication type by using the following command:

```
ip ospf authentication-type [message-digest | simple | none]
```

Variable definitions

The following table describes the command variables.

Variable	Value
`message-digest | simple | none`	Specifies the authentication type. • message-digest—MD5 digest authentication type • simple—simple password authentication type • none—no authentication type is applied to the interface

Configuring simple authentication keys for OSPF interfaces

Use this procedure to configure an interface authentication password.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure an interface authentication password by using the following command:

```
ip opsf authentication-key <password>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<password>	Specifies the password to be configured. This password can be up to 8 characters in length.

Defining MD5 keys for OSPF interfaces

Use this procedure to define the MD5 keys.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Define an MD5 key by using the following command:

```
ip ospf message-digest-key <key_number> md5 <key_value>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<key_number>	Specifies an index value for the MD5 key being configured. This is an integer value between 1 and 255.
<key_value>	Specifies the value of the MD5 key. This is a string value of up to 16 characters in length.

Displaying OSPF MD5 keys

Use this procedure to display OSPF MD5 key configuration.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF MD5 keys by using the following command:

```
show ip ospf authentication [interface vlan <vid>] [virtual-links]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[vlan <vid>]	Displays configured MD5 authentication keys for the specified interface. If no interface is specified, all interface MD5 keys are displayed.

Variable	Value
`virtual-links`	Displays configured MD5 authentication keys for virtual links.

Applying an MD5 key to an OSPF interface

Use this procedure to specify the primary MD5 key (configured using the **`ip ospf message-digest-key`** command) to use for authentication in instances where interface authentication uses an MD5 key.

Each OSPF interface supports up to 2 keys, identifiable by key ID, to facilitate a smooth key transition during the rollover process. Only the selected primary key is used to encrypt the OSPF transmit packets.

Assuming that all routers already use the same key for authentication and a new key is required, the process of key change is as follows:

1. Add the second key to all routers. The routers will continue to send OSPF packets encrypted with the old key.
2. Activate the second key on all routers by setting it as the primary key. Routers will send OSPF packets encrypted with the new key while still accepting packets using the old key. This is necessary as some routers will not have activated the new key.
3. After all routers activate the new key, remove the old key.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Apply the primary MD5 key by using the following command:

```
ip ospf primary-md5-key <key_id>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<key_id>	Specifies the index value for the MD5 key to apply. This is an integer value between 1 and 255.

Displaying OSPF interface authentication configuration

Use this procedure to display the authentication type and key applied to interfaces.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF authentication configuration for interfaces by using the following command:

```
show ip ospf int-auth
```

Configuring a virtual link

Use this procedure to create a virtual link.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Create a virtual interface by using the following command:

```
[default] [no] area virtual-link <area-id> <nghbr-router-id>
{[authentication-key <WORD>] [authentication-type {none|
simple|message-digest}] [primary-md5-key <1-255>] [dead-
```

```
interval <1-2147483647>] [hello-interval <1-65535>]
[retransmit-interval <1-3600>] [transit-delay <1-3600>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a virtual interface.
[default]	Configures the virtual link to default values.
<area_id>	Specifies the transit area ID in dotted decimal notation (A.B.C.D).
<nghbr-router-id>	Specifies the neighbor router ID expressed as an IP address.
authentication-key *<WORD>*	Specifies the unique identifier assigned to the authentication key.
authentication-type	Specifies one of the following authentication types: • none • simple • password • message digest MD5 TIP: Up to 2 MD5 keys are allowed for message digest. The default authentication type is none.
primary-md5-key	Specifies the user-selected key used to encrypt OSPF protocol packets for transmission.
dead-interval	Specifies the time interval, in seconds, that a Hello packet has not been transmitted from the virtual interface before its neighbors declare it down. Expressed as an integer from 1-2147483647, the default dead interval value is 60 seconds.
hello-interval	Specifies the time interval, in seconds, between transmission of Hello packets from the virtual interface. Expressed as an integer from 1-65535, the hello-interval default value is 10 seconds.
retransmit-interval	Specifies the time interval, in seconds, between link stage advertisement retransmissions for adjacencies belonging to the virtual interface. Expressed as an integer from 1-3600, the default value is 5 seconds.
transit-delay	Specifies the estimated number of seconds required to transmit a link state update packet over the virtual interface. Expressed as an integer from 1-3600, the default value is 1 second.

Creating a virtual interface message digest key

Use this procedure to create a virtual interface message digest key.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Create a virtual interface message digest key by using the following command:

`area virtual-link message-digest-key <area_id> <neighbor_id> <1-255> md5-key <WORD>`

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a virtual interface message digest key.
[default]	Specifies default values for the virtual interface message digest key.
<area_id>	Specifies the transit area Id expressed as an IP address.
<neighbor_id>	Specifies the neighbor router ID expressed as an IP address.
<1-255>	Specifies the primary MD5 key value, expressed as an integer from 1-255.
md5-key *<WORD>*	Specifies the user-selected key used to encrypt OSPF protocol packets for transmission.

Enabling automatic virtual links

Use this command to enable global automatic Virtual Link creation.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Enable global automatic Virtual Link creation by using the following command:

```
[default] [no] auto-vlink
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables global automatic Virtual Link creation.
[default]	Configures automatic Virtual Link creation to default.

Job aid: example of configuring automatic virtual links

Consider the following situation:

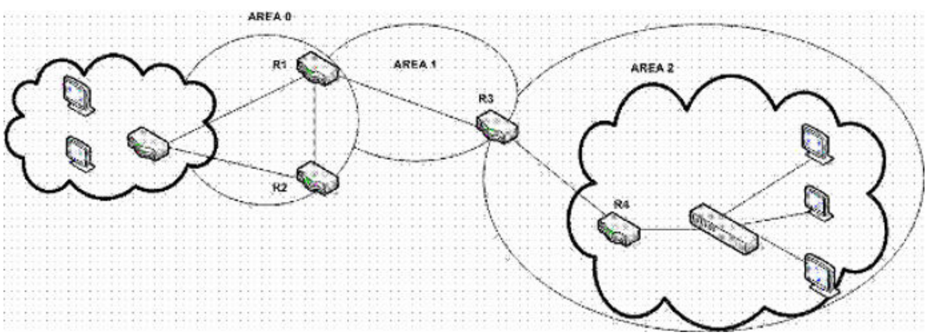

In this case, R4 in Area2 cannot be physically connected to Area0 (for some reason) and it will be connected to R3 which is NOT a backbone ABR (like R1 is for instance). As Area2 is not directly connected to backbone Area0 or directly connected to a backbone ABR router, clients from Area2 will not be able to access anything outside Area2. Also, router R3 is an ABR router connected to two non-backbone areas.

In order to solve these problems, virtual-link must be configured between router R3 and R1 which are both ABRs. Virtual-link cannot be configured on non-ABR routers.

Consider the following Router IDs:

- R1 : 1.0.0.0
- R3 : 3.0.1.0
- R4 : 4.0.2.0

The virtual-link can be configured in two ways on ABR routers :

- Configuring the virtual link manually
- Configuring the virtual link automatically

The following is an example for creating an auto virtual link:

Table 8: Creating auto virtual link

```
R1 (config-router)#auto-vlink
```

Example : 1

```
R1(config)#show ip ospf
Router ID:  1.0.0.0
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  67
New Link-State Advertisements Received:  722
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Enabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled

R3 (config-router)#auto-vlink
```

Example : 2

```
R3(config)#show ip ospf
Router ID:  3.0.1.0
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  67
New Link-State Advertisements Received:  722
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Enabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

The following is an example for deleting an auto virtual link:

Table 9: Deleting auto virtual link

```
R1 (config-router)#no auto-vlink
```

Example : 1

```
R1(config)#show ip ospf
Router ID:  1.0.0.0
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  67
New Link-State Advertisements Received:  722
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

```
R3 (config-router)#no auto-vlink
```

Example : 2

```
R3(config)#show ip ospf
Router ID:  3.0.1.0
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  67
New Link-State Advertisements Received:  722
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

Displaying OSPF virtual links

Use this procedure to display the configuration of OSPF virtual links.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF virtual links by using the following command:

```
show ip ospf virtual-links
```

Displaying OSPF virtual neighbors

Use this procedure to display OSPF virtual neighbor information.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF virtual neighbors by using the following command:

```
show ip ospf virtual-neighbors
```

Configuring an OSPF host route

Use this procedure to add a host to a router.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Add a host to a router by using the following command:

```
[no] host-route <A.B.C.D.> metric <0-65535>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a host route from the router.
<A.B.C.D.>	Specifies the host IP address.
[default]	Configures OSPF host route to default.
metric *<0-65535>*	Specifies an integer between 0 and 65535 representing the configured cost of the host route.

Job aid: example of configuring an OSPF host route

The following is an example for creating a host route:

```
R3(config)#router ospf R3(config-router)#host-route 11.11.11.111
metric 10 R3(config-router)#show ip ospf host-route
```

Host IP	Metric
11.11.11.111	10

Displaying OSPF host routes

Use this procedure to display OSPF host routes.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF host routes by using the following command:

```
show ip ospf host-route
```

Displaying the OSPF link state database

Use this procedure to display the OSPF link state database.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display the OSPF link state database by using the following command:

```
show ip ospf lsdb [adv-rtr <router_id>] [area <area-id>]
[detail <router_id>] [lsa-type <type>] [lsid <ip_address>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[adv-rtr *<router_id>*]	Displays OSPF LSDB information related to the specified advertisement router.
[area *<area-id>*]	Displays OSPF LSDB information related to the specified area.

Variable	Value
detail [<*router_id*>]	Display detailed OSPF LSDB information related to the specified advertisement router. If no router is specified, all detailed LSDB information is displayed.
[lsa-type <*type*>]	Displays OSPF LSDB information for the specified LSA type.
[lsid <*ip_address*>]	Displays OSPF LSDB information for the specified link state ID.

Displaying the external link state database

Use this procedure to display the external link state database.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF ASE LSAs by using the following command:

```
show ip ospf ase
```

Initiating an SPF run to update the OSPF LSDB

Manually initiate an SPF run to immediately update the link state database. Use this procedure, in the following situations:

- when you need to immediately restore a deleted OSPF-learned route
- as a debug mechanism when the routing table entries and the link-state database are not synchronized

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Immediately initiate an SPF run by using the following command:

```
ip ospf spf-run
```

Displaying OSPF default port metrics

Use this procedure to display OSPF default metrics for different port types.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF default metrics by using the following command:

```
show ip ospf default-cost
```

Displaying OSPF statistics

Use this procedure to display OSPF statistics.

Prerequisites

- Clear OSPF statistics counters by using the **`clear ip ospf counters`** command.
- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF statistics by using the following command:

```
show ip ospf stats
```

Displaying OSPF interface statistics

Use this procedure to display OSPF interface statistics.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display OSPF interface statistics by using the following command:

show ip ospf ifstats *<if-ip>* [mismatch] [detail]

Variable definitions

The following table describes the command variables.

Variable	Value
<if-ip>	Displays OSPF statistics for the specified interface IP address. Omitting this parameter displays statistics for the backbone area.
mismatch	Displays statistics where the area ID not matched.
detail	Display detailed statistics.

Clearing OSPF statistics counters

Use this procedure to clear OSPF statistics counters, including mismatch counters.

This procedure is applicable only to the base unit in a stack.

Prerequisites

• Log on to the Global Configuration mode in ACLI.

Procedure steps

Clear OSPF statistics counters by using the following command:

`clear ip ospf counters <1-4094>`

Variable definitions

The following table describes the command variables.

Variable	Value
<1-4094>	Specifies the VLAN ID. Range is 1-4094. If no VLAN is specified, the command clears OSPF global counters.

Chapter 8: OSPF configuration examples using ACLI

The following sections provide OSPF configuration examples using ACLI.

Basic OSPF configuration examples

This section contains examples to help you configure OSPF on your switch or stack. More advanced configuration examples can be found in Advanced OSPF configuration examples on page 133.

Note:

In many of the following configuration examples, a brouter port is used to create a connection to the network core. The use of a brouter port is only one of many ways to create such a connection.

Basic OSPF configuration

A basic OSPF configuration will learn OSPF routes from other OSPF devices and propagate routes to other OSPF devices. The following procedure describes the creation of a basic OSPF configuration:

1. Log into User EXEC mode.

   ```
   ERS4000>enable
   ```

2. Log into Global Configuration mode.

   ```
   ERS4000#config terminal
   ```

3. Enable IP routing globally.

   ```
   ERS4000(config)#ip routing
   ```

4. Enable OSPF globally.

   ```
   ERS4000(config)#router ospf enable
   ```

5. Log into the OSPF router configuration mode. It is not necessary to make any changes at this time but entering the router configuration mode is a good way to verify that the mode has been activated.

```
ERS4000(config)#router ospf
```

> **Important:**
>
> The remainder of this procedure refers to VLAN 35. Although VLAN 35 is used for this example, any port type VLAN could be used.

6. Return to Global Configuration mode.

   ```
   ERS4000(config-router)#exit
   ```

7. Create a port type VLAN as VLAN number 35 in spanning tree protocol group 1.

   ```
   ERS4000(config)#vlan create 35 type port 1
   ```

8. Log into the Interface Configuration mode for VLAN 35.

   ```
   ERS4000(config)#interface vlan 35
   ```

9. Enable IP routing on VLAN 35.

   ```
   ERS4000(config-if)#ip routing
   ```

10. Assign an IP address to VLAN 35.

    ```
    ERS4000(config-if)#ip address 1.1.2.25 255.255.255.0
    ```

11. Enable OSPF in VLAN 35.

    ```
    ERS4000(config-if)#ip ospf enable
    ```

12. Return to Global Configuration mode.

    ```
    ERS4000(config-if)#exit
    ```

13. By default all ports belong to a newly created VLAN. This command removes all of the ports from VLAN 35.

    ```
    ERS4000(config)#vlan members remove 35 all
    ```

14. Add ports 1 through 10 to VLAN 35.

    ```
    ERS4000(config)#vlan members add 35 1-10
    ```

Basic ASBR configuration

The Autonomous System Boundary Router (ASBR) is used in OSPF to import routes that come from non-OSPF sources such as:

- Local interfaces that are not part of OSPF.
- RIP interfaces.
- RIP learned routes.
- Static routes.

This quick reference will help in the configuration of OSPF to import these types of routes. This will allow the rest of the OSPF network to learn them as OSPF routes. To create a basic ASBR configuration, follow this procedure:

1. Log into User EXEC mode.

   ```
   ERS4000>enable
   ```

2. Log into Global Configuration mode.

   ```
   ERS4000#config terminal
   ```

3. Log into the OSPF router configuration mode.

   ```
   ERS4000(config)#router ospf
   ```

4. Enable ASBR functionality.

   ```
   ERS4000(config-router)#as-boundary-router enable
   ```

5. Use the following commands to select the type of routes that OSPF will distribute to other OSPF devices. RIP, direct, and static routes are supported.

   ```
   ERS4000(config-router)#redistribute rip enable
   ERS4000(config-router)#redistribute direct enable
   ERS4000(config-router)#redistribute static enable
   ```

6. Return to Global Configuration mode.

   ```
   ERS4000(config-router)#exit
   ```

7. Once the commands in step 5 have be used to select the types of routes to redistribute, apply the changes globally with the following commands.

   ```
   ERS4000(config)#ip ospf apply edistribute rip
   ERS4000(config)#ip ospf apply redistribute direct
   ERS4000(config)#ip ospf apply redistribute static
   ```

Configuring ECMP for OSPF using ACLI

For information about configuring ECMP for OSPF using ACLI, see the following sections:

- Setting the number of ECMP paths using ACLI on page 131
- Displaying the ECMP configuration using ACLI on page 132

Setting the number of ECMP paths using ACLI

Use this procedure to configure Equal Cost Multi Path (ECMP) for Open Shortest Path First (OSPF). You can specify up to four paths.

Prerequisites

- Enable routing on the switch
- Enable OSPF
- Use this command in the Global Configuration mode.

Procedure steps

Enter the following command:

```
ospf maximum-path <1-4>
```

Example

In the following example, you configure the router to use up to two equal-cost paths to reach any OSPF network destination.

```
ERS4000(config)# ospf maximum-path 2
```

Variable definitions

The following table describes the parameters for the **ospf maximum-path** command.

Variable definition

Variable	Value
<1–4>	Specifies the number of ECMP paths to use with OSPF in a range from 1 to 4. DEFAULT: 1

Displaying the ECMP configuration using ACLI

Use this procedure to display or verify the ECMP configuration on your switch.

Prerequisites

- Use this command in the User EXEC mode.

Procedure steps

Enter the following command:

```
show ecmp
```

Example

```
ERS4000(config)# show ecmp
Protocol   MAX-PATH
--------  ----------
static:        1
```

```
rip:            1
ospf:           4
```

Advanced OSPF configuration examples

This section contains examples of common OSPF-related configuration tasks.

The Avaya Ethernet Routing Switch 4000 Series supports the following OSPF standards:

- RFC 2328 (OSPF version 2)
- RFC 1850 (OSPF Management Information Base)
- RFC 2178 (OSPF MD5 cryptographic authentication)

This section provides examples of the common OSPF configuration tasks and includes the ACLI commands used to create the configuration.

Configuring an IP OSPF interface

You can configure an OSPF interface on a brouter port or on a VLAN. The following section demonstrates the creation of the example OSPF interface illustrated below.

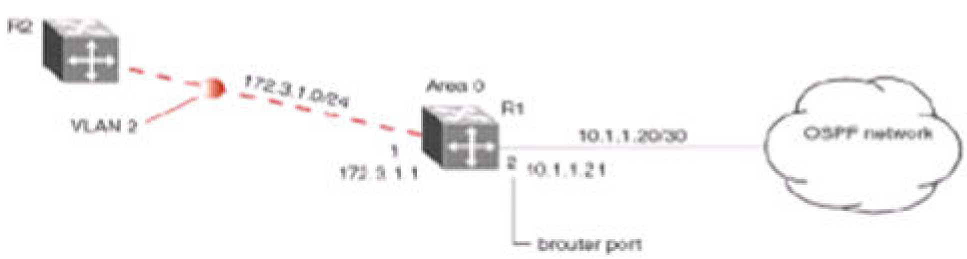

Figure 25: OSPF interface example topology

To create the OSPF interface illustrated in the preceding figure for router R1, follow this procedure:

1. Configure brouter port OSPF interface.

   Configure port 2 as a brouter port with VLAN ID of 2134 and enable OSPF on this interface

   Example

```
ERS4000# config terminal
ERS4000(config)# interface fast 2
```

```
ERS4000(config-if)# brouter port 2 vlan 2134 subnet 10.1.1.21/30
ERS4000(config-if)# router ospf
ERS4000(config-router)# network 10.1.1.21
```

2. Configure the VLAN OSPF interface.

   Create a port-based VLAN (VLAN 2) using spanning tree group 1, assign IP address 172.3.1.1 to VLAN 2 and enable OSPF on this interface.

   Example

```
ERS4000(config)# vlan create 2 type port
ERS4000(config)# spanning-tree stp 1 add-vlan 2
ERS4000(config)# vlan member add 2 1
ERS4000(config)# interface vlan 2
ERS4000(config-if)# ip address 172.3.1.1 255.255.255.0
ERS4000(config-if)# router ospf
ERS4000(config-router)# network 172.3.1.1
```

3. Assign a router ID to the new interface and enable OSPF globally.

   Example

```
ERS4000(config)# router ospf
ERS4000(config-router)# router-id 1.1.1.1
ERS4000(config-router)# exit
ERS4000(config)# router ospf enable
```

OSPF security configuration example using Message Digest 5

In the configuration example illustrated below, MD5 is configured between router R1 and R2.

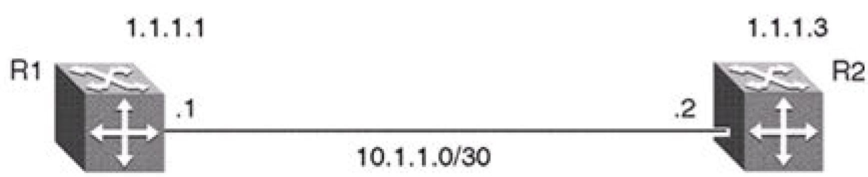

Figure 26: MD5 configuration example

To replicate the preceding configuration example using the key ID 2 and key value **qwsdf89**, perform the following steps:

1. Configure MD5 authentication on R1.

```
ERS4000(config)#interface vlan 2
ERS4000(config-if)#ip ospf message-digest-key 2 md5 qwsdf89
ERS4000(config-if)#ip ospf primary-md5-key 2
ERS4000(config-if)#ip ospf authentication-type message-digest
```

2. Configure MD5 authentication on R2.

```
ERS4000(config)#interface vlan 2
ERS4000(config-if)#ip ospf message-digest-key 2 md5 qwsdf89
```

```
ERS4000(config-if)#ip ospf primary-md5-key 2
ERS4000(config-if)#ip ospf authentication-type message-digest
```

Configuring OSPF network types

OSPF network types were created to allow OSPF-neighboring between routers over different types of network infrastructures. With this feature, each interface can be configured to support the various network types.

In the example configuration illustrated below, VLAN 2 on Avaya Ethernet Routing Switch 4000 Series R1 is configured for OSPF with the interface type field value set as **passive**. Because VLAN 2 is set as **passive**, OSPF hello messages are not sent on this segment, although R1 continues to advertise this interface to the remaining OSPF network.

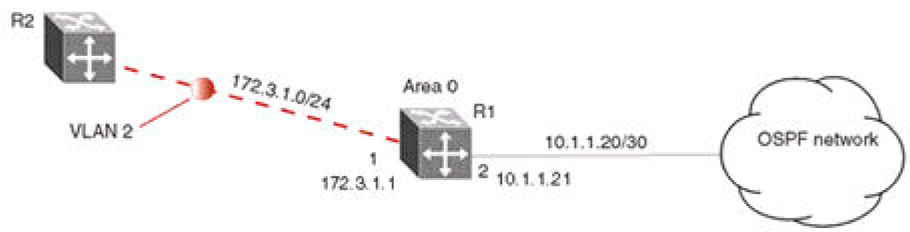

Figure 27: OSPF network example

To create the configuration illustrated in the preceding figure for router R1, use the following commands:

```
ERS4000(config)# vlan create 2 type port
ERS4000(config)# vlan mem add 2 1
ERS4000(config)# interface vlan 2
ERS4000(config-if)# ip address 172.3.1.1 255.255.255.0
ERS4000(config-if)# ip ospf network passive
```

The Avaya Ethernet Routing Switch 4000 Series supports the following types of networks:

- **Broadcast** - Automatically discovers every OSPF router on the network by sending OSPF hellos to the multicast group **AllSPFRouters** (224.0.0.5). Neighboring is automatic and requires no configuration. This interface type is typically used in an Ethernet environment.
- **Passive** - Allows interface network to be included in OSPF without generating LSAs or forming adjacencies. Typically used on an access network. This also limits the amount of CPU cycles required to process the OSPF routing algorithm.

Configuring Area Border Routers (ABR)

Configuration of an OSPF ABR is an automatic process on the Avaya Ethernet Routing Switch 4000 Series ; no user intervention is required. The Avaya Ethernet Routing Switch 4000 Series automatically becomes an OSPF ABR when it has operational OSPF interfaces belonging to more than one area.

In the configuration example below, the Avaya Ethernet Routing Switch 4000 Series R1 is automatically configured as an OSPF ABR after it is configured with an OSPF interface for area 0.0.0.0 and 0.0.0.2.

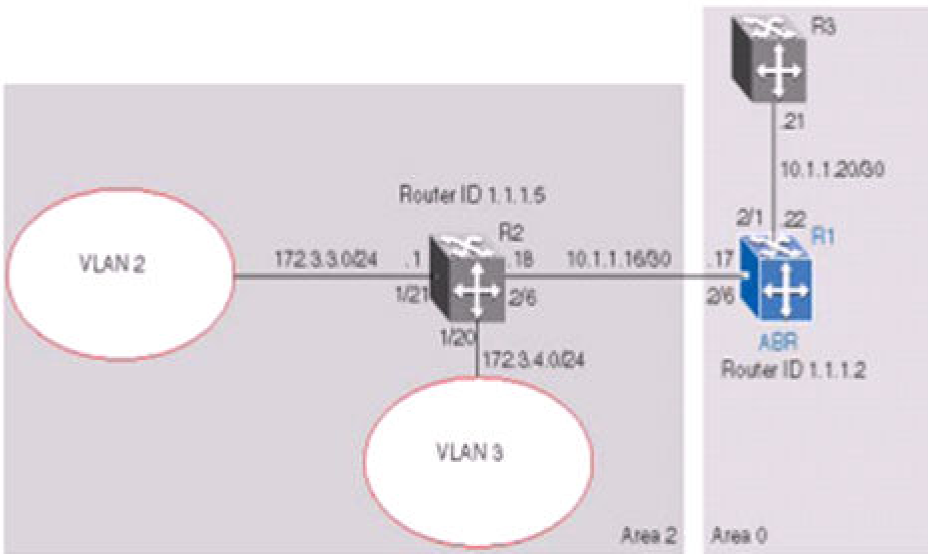

Figure 28: ABR configuration example

To recreate the illustrated ABR configuration, use the following procedure:

1. Configure an OSPF interface on port 2/6.

   Configure port 2/6 as a brouter port in VLAN 100.

   Example

   ```
   ERS4000(config)# interface fast 2/6
   ERS4000(config-if)# brouter port 2/6 vlan 100 subnet
   10.1.1.17/30
   ERS4000(config-if)#ip ospf enable area 0.0.0.2
   ```

2. Configure an OSPF interface on port 2/1.

   Configure port 2/1 as a brouter port in VLAN 200 and enable OSPF on this interface.

   Example

   ```
   ERS4000(config)# interface fast 2/1
   ERS4000(config-if)# brouter port 2/1 vlan 200 subnet
   10.1.1.22/30
   ERS4000(config-if)# ip ospf enable
   ```

3. Enable OSPF.

   Configure R1 as an ABR. Note that, by default, OSPF interface 10.1.1.22 is placed into OSPF area 0.0.0.0. Because one additional area of 0.0.0.2 is created and OSPF interface 10.1.1.17 is added to area 0.0.0.2, R1 automatically becomes an ABR.

   ```
   ERS4000(config-router)# router-id 1.1.1.2
   ERS4000(config-router)# area 0.0.0.2
   ERS4000(config-router)# network 10.1.1.17 area 0.0.0.2
   ERS4000(config)# router ospf enable
   ```

4. Configure area range.

   Configure R1 to enclose the two networks (172.3.3.0 and 172.3.4.0) into an address range entry 172.3.0.0 in area 0.0.0.2. R1 will generate a single summary advertisement into the backbone for 172.3.0.0 with metric 100.

   ```
   ERS4000(config-router)# area 0.0.0.2 range 172.3.0.0/16
   summary-link advertise-mode summarize advertise-metric 100
   ```

To display the created areas, use the **show ip ospf area** command. Usage of this command on the example configuration would yield the following output:

```
Area ID:  0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  2
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID:  0.0.0.2
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  2
Reachable Area Border Routers:  1
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
```

To display area ranges, use the **show ip ospf area-range** command. Usage of this command on the example configuration would yield the following output:

```
Area ID Range Subnet/Mask     Range Type                  Advertise Mode Metric
-------------    -----------    --    -----------    -------    ---------------------
0.0.0.2           172.3.0.0/16     Summary Link Summarize     100
```

To display ABR status, use the **show ip ospf** command. Usage of this command on the example configuration would yield the following output:

```
Router ID:  1.1.1.2
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  45698(0xb282)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  5
New Link-State Advertisements Received:  34
```

```
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

Configuring Autonomous System Border Routers (ASBR)

An ASBR is a router that has a connection to another Autonomous System to distribute any external routes that originated from a protocol into OSPF. An Avaya Ethernet Routing Switch 4000 Series configured as an ASBR can:

- Distribute all OSPF routes to RIP.
- Distribute RIP, direct, or static routes to OSPF.

Distributing OSPF routes to RIP and RIP to OSPF using AS-external LSA Type 1 metrics

The following configuration example displays an Avaya Ethernet Routing Switch 4000 Series configured as an ASBR between an OSPF and RIP version 2 network. In this example, the router distributes all OSPF routes to the RIP network and all RIP routes to the OSPF network.

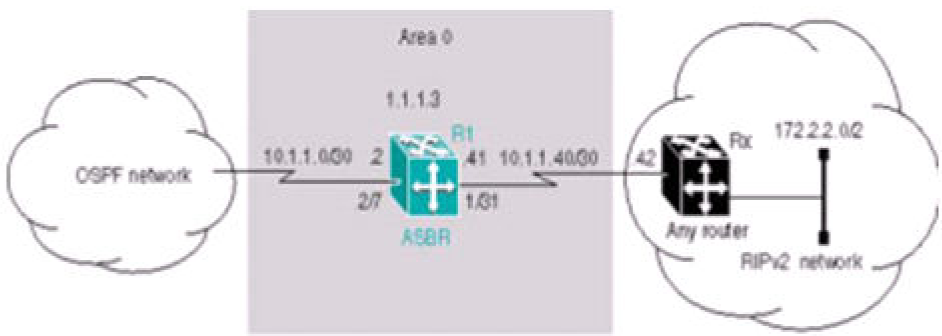

Figure 29: ASBR distribution example

Use the following procedure to replicate the ASBR distribution example:

1. Configure RIP.

   Configure the RIP interface on R1 by configuring port 1/31 as a brouter port in VLAN 100 and enabling RIP on this interface.

   Example

   ```
   ERS4000(config)# interface fast 1/31
   ERS4000(config-if)# brouter port 1/31 vlan 100 subnet 10.1.1.41/30
   ERS4000(config)# router rip
   ERS4000(config-router)# network 10.1.1.41
   ```

2. Configure the RIP interface for RIP version 2 mode only.

Example

```
ERS4000 (config)# router rip enable
ERS4000 (config)# interface vlan 100
ERS4000 (config-if)# ip rip receive version rip2 send version rip2
```

3. Configure the OSPF interface.

   Configure port 2/7 as a brouter port in VLAN 200 and enable OSPF on this interface.

   Example

```
ERS4000(config)# interface fast 2/7
ERS4000(config-if)# brouter port 2/7 vlan 200 subnet 10.1.1.2/30
ERS4000(config-if)# router ospf
ERS4000(config-router)# network 10.1.1.2
```

4. Make R1 the ASBR.

   Configure R1 as an ASBR and assign the OSPF Router-ID.

   Example

```
ERS4000(config)# router ospf
ERS4000(config-router)# as-boundary-router enable
ERS4000(config-router)# router-id 1.1.1.3
ERS4000(config)# router ospf enable
```

5. Configure OSPF route distribution.

   Example

   Configure OSPF route distribution to import RIP into OSPF. The Avaya Ethernet Routing Switch 4000 Series distributes the RIP routes as AS-external LSA (LSA type 5), using external metric type 1.

   Example

```
ERS4000(config)# router ospf
ERS4000(config-router)# redistribute rip enable metric 10 metric-type
type1
ERS4000(config)# ip ospf apply redistribute rip
```

6. Configure a route policy.

   A route policy is required for OSPF to RIP route redistribution. After you create the route policy, apply it to the RIP interface.

   The following command creates a route policy named **`allow`** which distributes both direct and OSPF interfaces.

   Example

```
ERS4000(config)# route-map allow permit 1 enable match protocol
direct,ospf
```

7. Apply the route policy to the RIP Out Policy.

The following commands apply the route policy to RIP interface 10.1.1.41.

```
ERS4000(config)# interface vlan 100
ERS4000(config-if)# ip rip out-policy allow
```

The configuration steps described in the preceding example distributes all OSPF routes to RIP. However, there are times when it can be more advantageous to distribute only a default route to RIP. The following configuration steps describe how to distribute only a default route to RIP instead of all OSPF routes to RIP.

To configure R1 to distribute a default route only to RIP, complete the following steps:

1. Configure an IP prefix list with a default route.

   The following command creates an IP prefix list named **default** with an IP address of 0.0.0.0.

   ```
   ERS4000(config)# ip prefix-list default 0.0.0.0/0
   ```

2. Configure a route policy.

   Create a route policy named **Policy_Default** which distributes the IP prefix list created in step 1. Note that **ospf** is selected as the **match-protocol** value. This causes the default route to be advertised through RIP only if OSPF is operational.

   ```
   ERS4000(config)# route-map Policy_Default permit 1 enable match protocol
   ospf set injectlist default
   ERS4000(config)# route-map Policy_Default 1 set  metric-type type1
   ```

3. Apply the route policy to the RIP Out Policy.

   Apply the route policy created in step 2 to RIP interface 10.1.1.41.

   ```
   ERS4000(config)# interface vlan 100
   ERS4000(config-if)# ip rip out-policy Policy_Default
   ```

Stub area configuration example

In the configuration example illustrated below, the Avaya Ethernet Routing Switch 4000 Series R1 is configured in Stub Area 2, and R2 is configured as a Stub ABR for Area 2.

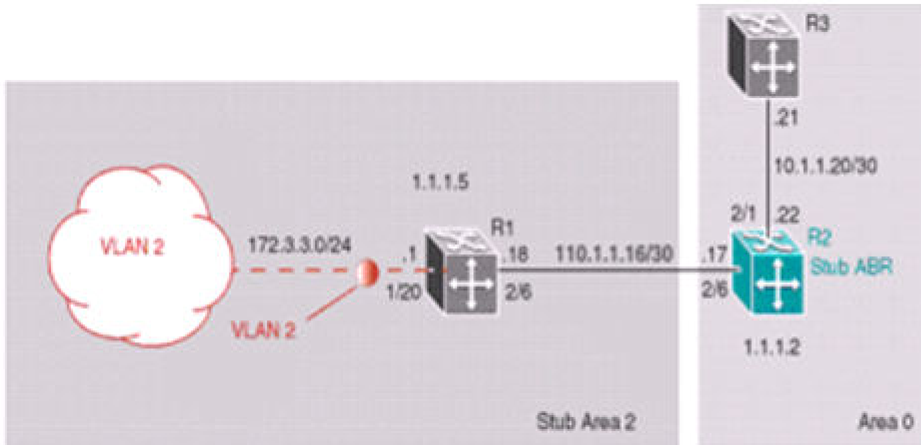

Figure 30: OSPF stub area example

Note:

AS-external LSAs are not flooded into a stub area. Instead, only one default route to external destinations is distributed into the stub area by the stub ABR router. The area default cost specifies the cost for advertising the default route into stub area by the ABR.

Use the following outlined procedure to perform the preceding stub area configuration illustration:

1. Configure router R1.

   Configure the OSPF interface on R1, configure port 2/6 as a brouter port in VLAN 100.

   Example

```
ERS4000(config)# interface fast 2/6
ERS4000(config-if)# brouter vlan 100 subnet 10.1.1.18/30
```

2. Configure VLAN 2 on R1.

   Create VLAN 2 and assign an IP address to it. Example

   Example

```
ERS4000(config)# vlan create 2 type port
ERS4000(config)# vlan mem add 2 1/20
ERS4000(config)# interface vlan 2
ERS4000(config-if)# ip address 172.3.3.1 255.255.255.0
```

3. Enable OSPF on R1.

   Configure R1 in stub area 2 with the Router-ID 1.1.1.5., add the OSPF interfaces to area 2 and enable OSPF on these interfaces.

   Example

```
ERS4000(config-router)# router-id 1.1.1.5
ERS4000(config-router)# area 0.0.0.2 import noexternal
```

```
ERS4000(config-router)# network 10.1.1.18 area 0.0.0.2
ERS4000(config-router)# network 172.3.3.1 area 0.0.0.2
ERS4000(config)# router ospf enable
```

4. Configure router R2.

   Configure the OSPF interface on R2, configure port 2/6 as a brouter port in VLAN 100.

   Example

```
ERS4000(config)# interface fast 2/6
ERS4000(config-if)# brouter port 2/6 vlan 100 subnet 10.1.1.17/30
```

5. Configure the second OSPF interface on R2.

   Configure port 2/1 as a brouter port in VLAN 300. Enable OSPF on this interface.

   Example

```
ERS4000(config)# interface fast 2/1
ERS4000(config-if)# brouter port 2/1 vlan 300 subnet 10.1.1.22/30
ERS4000(config-if)# ip ospf enable
```

6. Enable OSPF on R2.

   Configure R2 in stub area 2 with an area default cost of 10, disable import summary to prevent R2 from sending summary LSAs of area 0 into area 2 because R2 will originate only summary LSA for default route into area 2.

   Note:

   By default, OSPF interface 10.1.1.22 is placed into OSPF area 0.0.0.0. Because one additional area of 0.0.0.2 is added and OSPF interface 10.1.1.17 is added to area 0.0.0.2, R2 automatically becomes a stub ABR.

   Example

```
ERS4000(config-router)# router-id 1.1.1.2
ERS4000(config-router)# area 0.0.0.2 import noexternal
ERS4000(config-router)# no area 0.0.0.2 import-summary enable
ERS4000(config-router)# area 0.0.0.2 default-cost 10
ERS4000(config-router)# network 10.1.1.17 area 0.0.0.2
ERS4000(config)# router ospf enable
```

NSSA configuration example

The NSSA configuration example illustrated below demonstrates an Avaya Ethernet Routing Switch 4000 Series configured as a NSSA ASBR router.

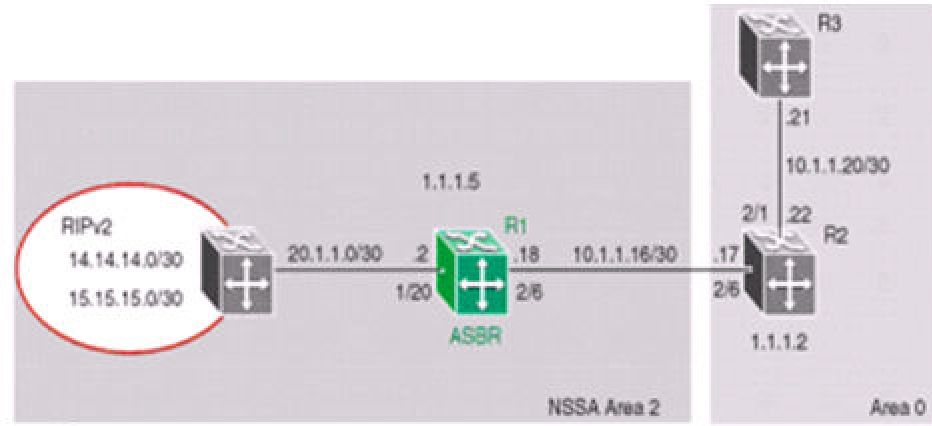

Figure 31: NSSA configuration example

To configure an NSSA, use the following procedure:

1. Configure router R1.

   Configure the RIP interface on R1 by configuring port 1/20 as VLAN 100 and enabling RIP on this interface.
2. Configure port 1/20 as a brouter port in VLAN 100 and enable RIP on this interface.

   Example

```
ERS4000(config)# interface fast 1/20
ERS4000(config-if)# brouter port 1/20 vlan 100 subnet 20.1.1.2/30
ERS4000(config)# router rip
ERS4000(config-router)# network 20.1.1.2
```
3. Enable RIP globally and configure the RIP version 2 interface.

   Example

```
ERS4000(config)# router rip enable
ERS4000(config-if)# ip rip receive version rip2 send version rip2
```
4. Configure the OSPF interface on R1.

   Configure port 2/6 as a brouter port in VLAN 200.

   Example

```
ERS4000(config)# interface fast 2/6
ERS4000(config-if)# brouter port 2/6 vlan 200 subnet 10.1.1.18/30
```
5. Enable OSPF on R1.

   Configure R1 as an ASBR, assign OSPF Router-ID 1.1.1.5, create OSPF NSSA area 2, add the OSPF interface 10.1.1.18 to area 2, and enable OSPF on the interface.

   Example

```
ERS4000 (config)# router ospf
ERS4000 (config-router)# as-boundary-router enable
```

```
ERS4000 (config-router)# router-id 1.1.1.5
ERS4000 (config-router)# area 0.0.0.2 import nssa
ERS4000 (config-router)# network 10.1.1.18 area 0.0.0.2
ERS4000 (config)# router ospf enable
```

6. Configure a route policy to distribute Direct and OSPF to RIP.

   Create a route policy named **`Rip_Dist`** that distributes directly connected and OSPF routes into RIP.

   Example

   ```
   ERS4000(config)# route-map Rip_Dist permit 1 enable match protocol
   direct,ospf set metric-type type1
   ```

7. Apply the **`Rip_Dist`** route policy to RIP Out Policy.

   Example

   ```
   ERS4000(config)# interface vlan 100
   ERS4000(config-if)# ip rip out-policy Rip_Dist
   ```

8. Configure OSPF route distribution to distribute RIP routes as AS-external LSA type 1.

   Example

   ```
   ERS4000(config)# router ospf
   ERS4000(config-router)# redistribute rip enable metric-type type1
   ERS4000(config)# ip ospf apply redistribute rip
   ```

Controlling NSSA external route advertisements

In an OSPF NSSA, the NSSA N/P-bit (in the OSPF hello packets Options field) is used to tell the ABR which external routes can be advertised to other areas. When the NSSA N/P-bit is set true, the ABR exports the external route. This is the default setting for the Avaya Ethernet Routing Switch 4000 Series . When the NSSA N/P-bit is not set true, the ABR drops the external route. A route policy can be created on the Avaya Ethernet Routing Switch 4000 Series to manipulate the N/ p-bit value.

For example, the illustration below shows a RIP network located in NSSA 2. If advertising the 15.15.15.0/24 network to area 0 is the only desired action, perform the following tasks:

- Enable R1 as an OSPF ASBR.
- Create NSSA area 0.0.0.2.
- Create a route policy to advertise OSPF and direct interfaces to RIP.
- Create a route policy to only advertise RIP network 15.15.15.0/24 to area 0 by using the NSSA N/P-bit.

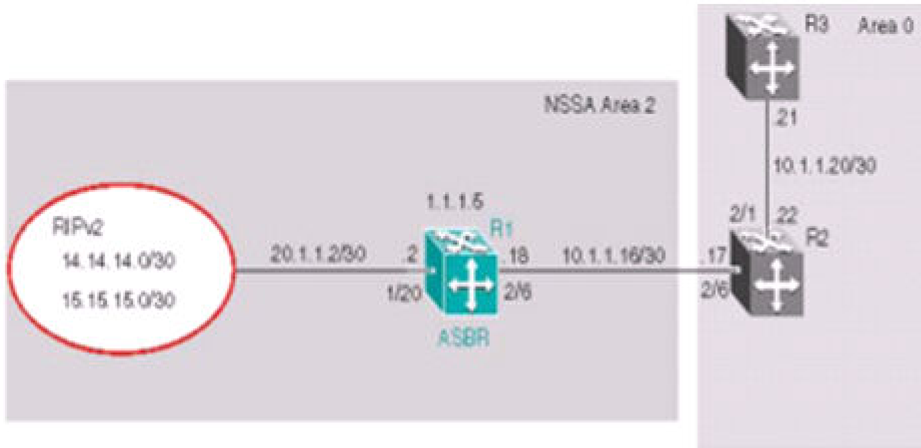

Figure 32: External route advertisement example

To configure an external route advertisement, use the following procedure:

1. Configure the RIP interface.

   Configure port 1/20 as a brouter port in VLAN 200 and enable RIP on this interface.

   Example:

```
ERS4000(config)# interface fast 1/20
ERS4000(config-if)# brouter port 1/20 vlan 200 subnet 20.1.1.2/30
ERS4000(config)# router rip
ERS4000(config-router)# network 20.1.1.2
```

2. Enable RIP globally and configure the RIP version 2 interface.

   Example

```
ERS4000(config)# router rip enable
ERS4000(config)# interface vlan 200
ERS4000(config-if)# ip rip receive version rip2 send version rip2
```

3. Configure the OSPF interface.

   Configure port 2/6 as a brouter port.

   Example

```
ERS4000(config)# interface fast 2/6
ERS4000(config-if)# brouter port 2/6 vlan 100 subnet 10.1.1.18/30
```

4. Enable OSPF.

   Configure R1 as an ASBR, assign the OSPF Router-ID 1.1.1.5, create OSPF NSSA area 2, add the OSPF interface 10.1.1.18 to area 2, and enable OSPF on the interface. Enable ASBR and OSPF globally.

   Example

```
ERS4000(config)#router ospf
ERS4000(config-router)#router-id 1.1.1.5
```

```
ERS4000(config-router)#as-boundary-router enable
ERS4000(config-router)#area 0.0.0.2 import nssa
ERS4000(config-router)#network 10.1.1.18 area 0.0.0.2
ERS4000(config)#router ospf enable
```

5. Create a route policy named **Rip_Dist** that distributes directly connected and OSPF routes into RIP.

   Example

   ```
   ERS4000(config)# route-map Rip_Dist permit 1 enable match protocol
   direct,ospf set metric-type type1
   ```

6. Apply route policy to RIP Out Policy.

   Example

   ```
   ERS4000(config)#interface vlan 200
   ERS4000(config-if)#ip rip out-policy Rip_Dist
   ```

7. Add two prefix lists (**15net** and **14net**) that are associated with the network addresses from the RIP version 2 network.

   Example

   ```
   ERS4000(config)#ip prefix-list 15net 15.15.15.0/24
   ERS4000(config)#ip prefix-list 14net 14.14.14.0/24
   ```

8. Create a route policy named **P_bit** that sets the NSSA N/P-bit only for the prefix list named **15net**.

   Example

   ```
   ERS4000(config)#route-map P_bit permit 1 enable match network 15net set
   nssa-pbit enable
   ERS4000(config)#route-map P_bit permit 2 enable match network 14net
   ERS4000(config)#no route-map P_bit 2 set nssa-pbit enable
   ```

9. Configure OSPF route distribution to distribute RIP routes as AS-external LSA Type 1.

   Example

   ```
   ERS4000(config)#router ospf
   ERS4000(config-router)#redistribute rip enable metric-type type1 route-
   policy P_bit
   ERS4000(config)#ip ospf apply redistribute rip
   ```

Configuring a multi-area complex

The multi-area complex configuration example described in this section uses five Avaya Ethernet Routing Switch 4000 Series devices (R1 to R5) in a multi-area configuration.

Many of the concepts and topology descriptions that are used in this example configuration are described in the previous sections of this chapter. The concepts shown in those examples are combined in this example configuration to show a real world topology example with command descriptions.

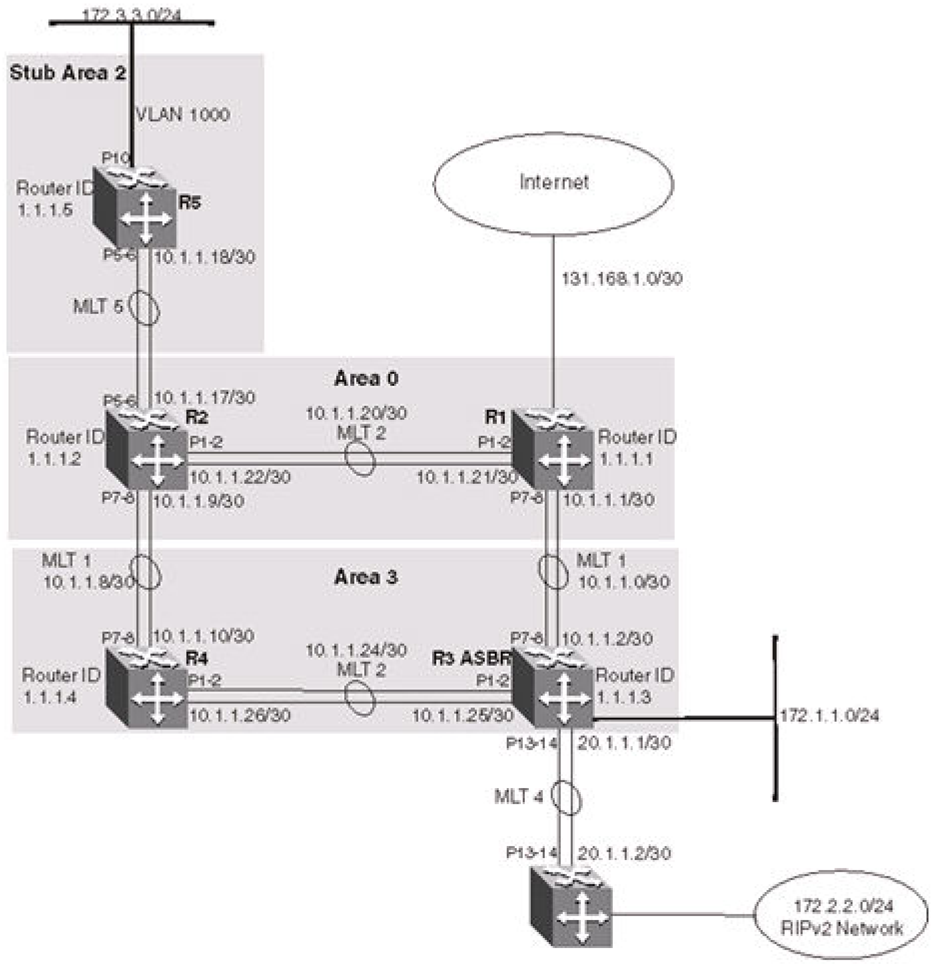

Figure 33: Multi-area complex example

For this configuration example, the Avaya Ethernet Routing Switch 4000 Series devices R1 through R5 are configured as follows:

- R1 is an OSPF ABR that is associated with OSPF Area 0 and 3.
- R2 is an OSPF Stub ABR for OSPF Area 2 and ABR to OSPF Area 3.
- R3 is an OSPF ASBR and is configured to distribute OSPF to RIP and RIP to OSPF.
- R4 is an OSPF internal router in Area 3.
- R5 is an internal OSPF stub router in Area 2.
- All interfaces used for this configuration are ethernet, therefore the OSPF interfaces are broadcast.
- The interface priority value on R5 is set to 0, therefore R5 cannot become a designated router (DR).
- Configure the OSPF Router Priority so that R1 becomes the DR (priority of 100) and R2 becomes backup designated router (BDR) with a priority value of 50.

Stub and NSSA areas are used to reduce the LSDB size by excluding external LSAs. The stub ABR advertises a default route into the stub area for all external routes.

The following list describes the commands used to create the illustrated configuration. A similar listing could be provided by using the **show running-config** command.

The following commands illustrate the status of the routers in the configuration example. Accompanying each command is the output matching to the configuration example.

R1 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree stp 2 create
spanning-tree stp 3 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 2 priority 8000
spanning-tree stp 2 hello-time 2
spanning-tree stp 2 max-age 20
spanning-tree stp 2 forward-time 15
spanning-tree stp 2 tagged-bpdu enable tagged-bpdu-vid 4002
spanning-tree stp 2 multicast-address 01:80:c2:00:00:00
spanning-tree stp 3 priority 8000
spanning-tree stp 3 hello-time 2
spanning-tree stp 3 max-age 20
spanning-tree stp 3 forward-time 15
spanning-tree stp 3 tagged-bpdu enable tagged-bpdu-vid 4003
spanning-tree stp 3 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol autopvid
auto-pvid
vlan name 1 "VLAN #1"
vlan create 102 name "VLAN #102" type port
vlan create 103 name "VLAN #103" type port
vlan ports 1-24 tagging unTagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan ports 25-26 tagging tagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan members 1 24-26 vlan members 102 1-2
vlan members 103 7-8
vlan ports 1-2 pvid 102
vlan ports 3-6 pvid 1
vlan ports 7-8 pvid 103
vlan ports 9-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 102 snooping disable
vlan igmp 102 proxy disable robust-value 2 query-interval 125
vlan igmp 103 snooping disable
vlan igmp 103 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
```

R1 configuration commands

```
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 2 add-vlan 102
spanning-tree stp 3 add-vlan 103
spanning-tree stp 2 enable
spanning-tree stp 3 enable interface Ethernet ALL
spanning-tree port 24-26 learning normal
spanning-tree port 1-2 stp 2 learning normal
spanning-tree port 7-8 stp 3 learning normal
spanning-tree port 24-26 cost 1 priority 80
spanning-tree port 1-2 stp 2 cost 1 priority 80
spanning-tree port 7-8 stp 3 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26
enable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 3 learning normal
mlt spanning-tree 2 stp 2 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing
interface vlan 102
ip address 10.1.1.21 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 103
ip address 10.1.1.1 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.1
no as-boundary-router enable
no trap enable timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
exit
enable
configure terminal
interface vlan 103
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 100
```

R1 configuration commands

`ip ospf authentication-type none`
`ip ospf mtu-ignore enable`
`no ip ospf advertise-when-down enable`
`ip ospf transmit-delay 1`
`ip ospf retransmit-interval 5`
`ip ospf hello-interval 10`
`ip ospf dead-interval 40`
`ip ospf enable`
`exit`
`interface vlan 102`
`ip ospf area 0.0.0.0`
`ip ospf network broadcast`
`ip ospf priority 100`
`ip ospf authentication-type none`
`ip ospf mtu-ignore enable`
`no ip ospf advertise-when-down enable`
`ip ospf enable`
`exit`

R2 configuration commands

`! *** STP (Phase 1) *** !`
`spanning-tree stp 2 create`
`spanning-tree cost-calc-mode dot1d`
`spanning-tree port-mode normal`
`spanning-tree stp 1 priority 8000`
`spanning-tree stp 1 hello-time 2`
`spanning-tree stp 1 max-age 20`
`spanning-tree stp 1 forward-time 15`
`spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001`
`spanning-tree stp 1 multicast-address 01:80:c2:00:00:00`
`spanning-tree stp 2 priority 8000`
`spanning-tree stp 2 hello-time 2`
`spanning-tree stp 2 max-age 20`
`spanning-tree stp 2 forward-time 15`
`spanning-tree stp 2 tagged-bpdu enable tagged-bpdu-vid 4002`
`spanning-tree stp 2 multicast-address 01:80:c2:00:00:00`
`! *** VLAN *** !`
`vlan configcontrol autopvid`
`auto-pvid`
`vlan name 1 "VLAN #1"`
`vlan create 100 name "VLAN #100" type port`
`vlan create 101 name "VLAN #101" type port`
`vlan create 102 name "VLAN #102" type port`
`vlan ports 1-2 tagging tagAll filter-untagged-frame disable filter-`
`unregistered-frames enable priority 0`
`vlan ports 3-6 tagging unTagAll filter-untagged-frame disable filter-`
`unregistered-frames enable priority 0`
`vlan ports 7-8 tagging tagAll filter-untagged-frame disable filter-`
`unregistered-frames enable priority 0`
`vlan ports 9-26 tagging unTagAll filter-untagged-frame disable filter-`
`unregistered-frames enable priority 0`
`vlan members 1 1-26`
`vlan members 100 5-6`
`vlan members 101 7-8`
`vlan members 102 1-2`
`vlan ports 1-2 pvid 102`
`vlan ports 3-4 pvid 1`
`vlan ports 5-6 pvid 100`

R2 configuration commands

```
vlan ports 7-8 pvid 101
vlan ports 9-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 100 snooping disable
vlan igmp 100 proxy disable robust-value 2 query-interval 125
vlan igmp 101 snooping disable
vlan igmp 101 proxy disable robust-value 2 query-interval 125
vlan igmp 102 snooping disable
vlan igmp 102 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
mlt 5 name "Trunk #5" enable member 5-6 learning normal
mlt 5 learning normal
mlt 5 bpdu all-ports
mlt 5 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 1 add-vlan 100
spanning-tree stp 2 add-vlan 101
spanning-tree stp 2 add-vlan 102
spanning-tree stp 2 enable
interface Ethernet ALL
spanning-tree port 1-26 learning normal
spanning-tree port 1-2,7-8 stp 2 learning normal
spanning-tree port 1-26 cost 1 priority 80
spanning-tree port 1-2,7-8 stp 2 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 1 learning normal
mlt spanning-tree 1 stp 2 learning normal
mlt spanning-tree 2 stp 1 learning normal
mlt spanning-tree 2 stp 2 learning normal
mlt spanning-tree 5 stp 1 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing
interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 100
ip address 10.1.1.17 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 101 ip address 10.1.1.9 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
```

R2 configuration commands

```
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 102 ip address 10.1.1.22 255.255.255.252 4
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** ECMP *** !
maximum-path 1 rip
maximum-path 1 ospf
maximum-path 1
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.2
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.2 import noexternal
default-cost 1
area 0.0.0.2 import-summaries enable
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
exit
enable
configure terminal
interface vlan 101
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 50
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable exit interface vlan 100
ip ospf area 0.0.0.2
ip ospf network broadcast
ip ospf priority 50
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 102
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 50
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
```

R2 configuration commands

```
ip ospf enable
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
```

R3 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree stp 3 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree nstp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 3 priority 8000
spanning-tree stp 3 hello-time 2
spanning-tree stp 3 max-age 20
spanning-tree stp 3 forward-time 15
spanning-tree stp 3 tagged-bpdu enable tagged-bpdu-vid 4003
spanning-tree stp 3 multicast-address 01:80:c2:00:00:00
! ***VLAN *** !
vlan configcontrol automatic
auto-pvid
vlan name 1 "VLAN #1"
vlan create 103 name "VLAN #103" type port
vlan create 104 name "VLAN #104" type port
vlan create 105 name "VLAN #105" type port
vlan create 1001 name "VLAN #1001" type port
vlan ports 1-2 tagging tagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan ports 3-6 tagging unTagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan ports 7-8 tagging tagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan ports 9-26 tagging unTagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan members 1 4-6,9,12,15-26
vlan members 103 7-8
vlan members 104 1-2
vlan members 105 13-14
vlan members 1001 10
vlan ports 1-2 pvid 104
vlan ports 3-6 pvid 1
vlan ports 7-8 pvid 103
vlan ports 9 pvid 1
vlan ports 10 pvid 1001
vlan ports 11-12 pvid 1
vlan ports 13-14 pvid 105
vlan ports 15-26 pvid 1
```

R3 configuration commands

```
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 103 snooping disable
vlan igmp 103 proxy disable robust-value 2 query-interval 125
vlan igmp 104 snooping disable
vlan igmp 104 proxy disable robust-value 2 query-interval 125
vlan igmp 105 snooping disable
vlan igmp 105 proxy disable robust-value 2 query-interval 125
vlan igmp 1001 snooping disable
vlan igmp 1001 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
mlt 4 name "Trunk #4" enable member 13-14 learning normal
mlt 4 learning normal
mlt 4 bpdu all-ports
mlt 4 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 3 add-vlan 103
spanning-tree stp 3 add-vlan 104
spanning-tree stp 1 add-vlan 105
spanning-tree stp 1 add-vlan 1001
spanning-tree stp 3 enable
interface Ethernet ALL
spanning-tree port 4-6,9,12-26 learning normal
spanning-tree port 1-2,7-8 stp 3 learning normal
spanning-tree port 4-6,9,12-26 cost 1 priority 80
spanning-tree port 1-2,7-8 stp 3 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
interface Ethernet ALL
spanning-tree port 10 learning disable
exit
interface Ethernet ALL exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 3 learning normal
mlt spanning-tree 2 stp 3 learning normal
mlt spanning-tree 4 stp 1 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing
interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 103
ip address 10.1.1.2 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
```

R3 configuration commands

```
exit
interface vlan 104 ip address 10.1.1.25 255.255.255.252 4
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 105
ip address 20.1.1.1 255.255.255.0 5
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 1001 ip address 172.1.1.1 255.255.255.0 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** Route Policies *** !
route-map Allow permit 1
route-map Allow 1 enable
route-map Allow 1 match protocol direct,ospf
no route-map Allow 1 match interface
route-map Allow 1 match metric 0
no route-map Allow 1 match network
no route-map Allow 1 match next-hop
route-map Allow 1 match
route-type any
no route-map Allow 1 match route-source
no route-map Allow 1 set injectlist
route-map Allow 1 set mask 0.0.0.0
route-map Allow 1 set metric 5
route-map Allow 1 set nssa-pbit enable
route-map Allow 1 set ip-preference 0
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.3
as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.0 import external
area 0.0.0.0 import-summaries enable
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
redistribute direct metric 10 metric-type type2 subnets allow
redistribute direct enable
redistribute rip metric 10 metric-type type2 subnets allow
redistribute rip enable
exit
enable
configure terminal
interface vlan 103
ip ospf area 0.0.0.3
ip ospf network broadcast
```

R3 configuration commands
ip ospf priority 1 ip ospf authentication-type none ip ospf mtu-ignore enable no ip ospf advertise-when-down enable ip ospf transmit-delay 1 ip ospf retransmit-interval 5 ip ospf hello-interval 10 ip ospf dead-interval 40 ip ospf enable exit interface vlan 104 ip ospf area 0.0.0.3 ip ospf network broadcast ip ospf priority 1 ip ospf authentication-type none ip ospf mtu-ignore enable no ip ospf advertise-when-down enable ip ospf enable exit interface vlan 105 ip ospf area 0.0.0.0 ip ospf network broadcast ip ospf priority 1 ip ospf authentication-type none ip ospf mtu-ignore enable no ip ospf advertise-when-down enable no ip ospf enable exit interface vlan 1001 ip ospf area 0.0.0.3 ip ospf network broadcast ip ospf priority 1 ip ospf authentication-type none ip ospf mtu-ignore enable no ip ospf advertise-when-down enable ip ospf enable exit interface vlan 1 ip ospf area 0.0.0.0 ip ospf network broadcast ip ospf priority 1 ip ospf authentication-type none ip ospf mtu-ignore enable no ip ospf advertise-when-down enable no ip ospf enable exit ! *** RIP *** ! router rip router rip enable timers basic holddown 120 timers basic timeout 180 update 30 default-metric 8 no network 10.1.1.2 no network 10.1.1.25 network 20.1.1.1 no network 172.1.1.1 no network 203.203.100.52 exit enable configure terminal interface vlan 103 no ip rip advertise-when-down enable no ip rip auto-aggregation enable no ip rip default-listen enable no ip rip default-supply enable ip rip cost 1 ip rip holddown

R3 configuration commands
120 ip rip listen enable no ip rip poison enable no ip rip proxy-announce enable ip rip receive version rip1OrRip2 ip rip send version rip1Comp ip rip timeout 180 no ip rip triggered enable ip rip supply enable exit interface vlan 104 no ip rip advertise-when-down enable no ip rip auto-aggregation enable no ip rip default-listen enable no ip rip default-supply enable ip rip cost 1 ip rip holddown 120 ip rip listen enable no ip rip poison enable no ip rip proxy-announce enable ip rip receive version rip1OrRip2 ip rip send version rip1Comp ip rip timeout 180 no ip rip triggered enable ip rip supply enable exit interface vlan 105 no ip rip advertise-when-down enable no ip rip auto-aggregation enable no ip rip default-listen enable no ip rip default-supply enable ip rip cost 1 ip rip holddown 120 ip rip listen enable ip rip out-policy Allow no ip rip poison enable no ip rip proxy-announce enable ip rip receive version rip1OrRip2 ip rip send version rip1Comp ip rip timeout 180 no ip rip triggered enable ip rip supply enable exit interface vlan 1001 no ip rip advertise-when-down enable no ip rip auto-aggregation enable no ip rip default-listen enable no ip rip default-supply enable ip rip cost 1 ip rip holddown 120 ip rip listen enable no ip rip poison enable no ip rip proxy-announce enable ip rip receive version rip1OrRip2 ip rip send version rip1Comp ip rip timeout 180 no ip rip triggered enable ip rip supply enable exit interface vlan 1 no ip rip advertise-when-down enable no ip rip auto-aggregation enable no ip rip default-listen enable

R3 configuration commands

```
no ip rip default-supply enable
ip rip cost 1
ip rip holddown 120
ip rip listen enable
no ip rip poison enable
no ip rip proxy-announce enable
ip rip receive version rip1OrRip2
ip rip send version rip1Comp
ip rip timeout 180
no ip rip triggered enable
ip rip supply enable
exit
```

R4 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree stp 3 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 3 priority 8000
spanning-tree stp 3 hello-time 2
spanning-tree stp 3 max-age 20
spanning-tree stp 3 forward-time 15
spanning-tree stp 3 tagged-bpdu enable tagged-bpdu-vid 4003
spanning-tree stp 3 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol automatic
auto-pvid
vlan name 1 "VLAN #1"
vlan create 101 name "VLAN #101" type port
vlan create 104 name "VLAN #104" type port
vlan ports 1-26 tagging unTagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan members 1 3-6,9-26
vlan members 101 7-8
vlan members 104 1-2
vlan ports 1-2 pvid 104
vlan ports 3-6 pvid 1
vlan ports 7-8 pvid 101
vlan ports 9-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 101 snooping disable
vlan igmp 101 proxy disable robust-value 2 query-interval 125
vlan igmp 104 snooping disable
vlan igmp 104 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
```

R4 configuration commands

```
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 3 add-vlan 101
spanning-tree stp 3 add-vlan 104
spanning-tree stp 3 enable interface Ethernet ALL
spanning-tree port 3-6,9-26 learning normal
spanning-tree port 1-2,7-8 stp 3 learning normal
spanning-tree port 3-6,9-26 cost 1 priority 80
spanning-tree port 1-2,7-8 stp 3 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 3 learning normal
mlt spanning-tree 2 stp 3 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 101
ip address 10.1.1.10 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay exit interface vlan 104
ip address 10.1.1.26 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.4
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.0 import external
area 0.0.0.0 import-summaries enable
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
exit
enable configure terminal
interface vlan 101
ip ospf  area 0.0.0.3
ip ospf network broadcast
ip ospf priority 1
```

R4 configuration commands

```
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 104
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
```

R5 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol autopvid
auto-pvid
vlan name 1 "VLAN #1"
vlan create 100 name "VLAN #100" type port
vlan create 1000 name "VLAN #1000" type port
vlan ports 1-26 tagging unTagAll filter-untagged-frame disable filter-
unregistered-frames enable priority 0
vlan members 1 24-26
vlan members 100 5-6
vlan members 1000 10
vlan ports 1-4 pvid 1
vlan ports 5-6 pvid 100
vlan ports 7-9 pvid 1
vlan ports 10 pvid 1000
vlan ports 11-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 100 snooping disable
```

R5 configuration commands

```
vlan igmp 100 proxy disable robust-value 2 query-interval 125
vlan igmp 1000 snooping disable
vlan igmp 1000 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
mlt 5 name "Trunk #5" enable member 5-6 learning normal
mlt 5 learning normal
mlt 5 bpdu all-ports
mlt 5 loadbalance basic
*** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 1 add-vlan 100
spanning-tree stp 1 add-vlan 1000
interface Ethernet ALL
spanning-tree port 5-6,24-26 earning normal
spanning-tree port 5-6,24-26 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
interface Ethernet ALL
spanning-tree port 10 learning disable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 5 stp 1 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 100
ip address 10.1.1. 18 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 1000
ip address 172.3.3.1 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp no
ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.5
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.0 import external
area 0.0.0.0 import-summaries enable
area 0.0.0.2 import noexternal
default-cost 1
```

R5 configuration commands

```
area 0.0.0.2 import-summaries enable
exit
enable
configure terminal
interface vlan 100
ip ospf area 0.0.0.2
ip ospf network broadcast
ip ospf priority 0
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 1000
ip ospf area 0.0.0.2
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
```

Router R1 Status

show vlan

```
Id Name   Type     Protocol User     PID        Active IVL/SVL Mgmt
----------------------------------------------------------
1   VLAN #1    Port None     0x0000    Yes    IVL     Yes
Port Members: 1-2,5-7,9-14,16-17,19-26  2

VLAN #2   Port None     0x0000   Yes    IVL     No
Port Members: 3-4,8,18  5

VLAN #5   Port None     0x0000   Yes    IVL     No
Port Members: 15
```

Total VLANs:3

show vlan ip

```
=========================================================================
Id  ifIndex Address        Mask          MacAddress       Offset Routing
=========================================================================
                            Primary Interfaces
-------------------------------------------------------------------------
-
1   10001   10.100.111.200 255.255.255.0  00:11:F9:35:84:40 1     Enabled
2   10002   3.3.3.1        255.255.255.0  00:11:F9:35:84:41 2     Enabled
5   10005   10.10.10.1     255.255.255.0  00:11:F9:35:84:44 5     Enabled
-------------------------------------------------------------------------
-
                            Secondary Interfaces
-------------------------------------------------------------------------
-
2   14096   4.4.4.1        255.255.255.0  00:11:F9:35:84:42 3     Enabled
2   18190   5.5.5.1        255.255.255.0  00:11:F9:35:84:43 4     Enabled
```

show ip ospf

```
Router ID:  1.1.1.1
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  49786(0xc27a)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  427
New Link-State Advertisements Received:  811
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID:  0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  35
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  15
Link-State Advertisements Checksum:  551120(0x868d0)
Area ID:  0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  37
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  454461(0x6ef3d)
```

show ip ospf interface

```
Interface:  10.1.1.1
Area ID:  0.0.0.3
```

show ip ospf interface

```
Admin State:  Enabled
Type:  Broadcast
Priority:  100
Designated Router:  10.1.1.1
Backup Designated Router:  10.1.1.2
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.21
Area ID:  0.0.0.0
Admin State:  Enabled
Type:  Broadcast
Priority:  100
Designated Router:  10.1.1.21
Backup Designated Router:  10.1.1.22
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf neighbor

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------- -------------- --- ----- ----------- ----
10.1.1.1  1.1.1.3       10.1.1.2       1   Full  0           Dyn
10.1.1.21 1.1.1.2       10.1.1.22      50  Full  0           Dyn
Total OSPF Neighbors:  2
```

show ip route

```
==================================================================
Ip Route
==================================================================
DST         MASK            NEXT      COST VLAN PORT PROT TYPE PRF
------------------------------------------------------------------
172.2.2.0   255.255.255.0   10.1.1.2  10   103  T#1  O    IB   120
172.1.1.0   255.255.255.0   10.1.1.2  20   103  T#1  O    IB    20
172.3.3.0   255.255.255.252 10.1.1.22 30   102  T#2  O    IB    25
20.1.1.0    255.255.255.0   10.1.1.2  10   103  T#1  O    IB   120
10.1.1.24   255.255.255.252 10.1.1.2  20   103  T#1  O    IB    20
10.1.1.20   255.255.255.252 10.1.1.21  1   102  ---- C    DB     0
10.1.1.16   255.255.255.252 10.1.1.22 20   102  T#2  O    IB    25
10.1.1.0    255.255.255.252 10.1.1.1   1   103  ---- C    DB     0
10.1.1.8    255.255.255.252 10.1.1.2  30   103  T#1  O    IB    20
Total Routes: 9
-----------------------------------------------------------------
TYPE Legend: I=Indirect Route, D=Direct Route, A=Alternative Route,
B=Best Route, E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R2 Status

show vlan

```
Id  Name      Type Protocol User PID Active IVL/SVL Mgmt
--- --------- ---- -------- -------- ------ ------- ----
1   VLAN #1   Port None     0x0000   Yes    IVL     Yes
Port Members: 1-26
100 VLAN #100 Port None     0x0000   Yes    IVL     No
Port Members: 5-6
101 VLAN #101 Port None     0x0000   Yes    IVL     No
Port Members: 7-8
102 VLAN #102 Port None     0x0000   Yes    IVL     No
Port Members: 1-2
```

show vlan ip

```
Id ifIndex Address        Mask            MacAddress        Offset Routing
1  10001   203.203.100.53 255.255.255.0   00:15:9B:F3:70:40 1      Enabled
100 10100  10.1.1.17      255.255.255.252 00:15:9B:F3:70:41 2      Enabled
101 10101  10.1.1.9       255.255.255.252 00:15:9B:F3:70:42 3      Enabled
102 10102  10.1.1.22      255.255.255.252 00:15:9B:F3:70:43 4      Enabled
```

show ip ospf

```
Router ID:  1.1.1.2
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  49786(0xc27a)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  99
New Link-State Advertisements Received:  66
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID:  0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  8
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  15
Link-State Advertisements Checksum:  551120(0x868d0)
Area ID:  0.0.0.2
Import Summaries:  Yes
Import Type:  No
External Intra-Area SPF Runs:  10
Reachable Area Border Routers:  1
Reachable Autonomous System Border Routers:  0
```

show ip ospf area
Link-State Advertisements: 9 Link-State Advertisements Checksum: 274851(0x431a3) Stub Metric: 1 Stub Metric Type: OSPF Metric Area ID: 0.0.0.3 Import Summaries: Yes Import Type: External Intra-Area SPF Runs: 13 Reachable Area Border Routers: 2 Reachable Autonomous System Border Routers: 1 Link-State Advertisements: 13 Link-State Advertisements Checksum: 454461(0x6ef3d)

show ip ospf interface
Interface: 10.1.1.9 Area ID: 0.0.0.3 Admin State: Enabled Type: Broadcast Priority: 50 Designated Router: 10.1.1.9 Backup Designated Router: 10.1.1.10 Authentication Type: None MTU Ignore: Yes Advertise When Down: No Metric Value: 10 Interface: 10.1.1.17 Area ID: 0.0.0.2 Admin State: Enabled Type: Broadcast Priority: 50 Designated Router: 10.1.1.17 Backup Designated Router: 0.0.0.0 Authentication Type: None MTU Ignore: Yes Advertise When Down: No Metric Value: 10 Interface: 10.1.1.22 Area ID: 0.0.0.0 Admin State: Enabled Type: Broadcast Priority: 50 Designated Router: 10.1.1.21 Backup Designated Router: 10.1.1.22 Authentication Type: None MTU Ignore: Yes Advertise When Down: No Metric Value: 10 Interface: 203.203.100.53 Area ID: 0.0.0.0 Admin State: Disabled Type: Broadcast Priority: 1 Designated Router: 0.0.0.0 Backup Designated Router: 0.0.0.0 Authentication Type: None MTU Ignore: Yes Advertise When Down: No Metric Value: 10

show ip ospf neighbor

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------- -------------- --- ----- ----------- ----
10.1.1.9  1.1.1.4       10.1.1.10      1   Full  0           Dyn
10.1.1.17 1.1.1.5       10.1.1.18      0   Full  0           Dyn
10.1.1.22 1.1.1.1       10.1.1.21      100 Full  0           Dyn
Total OSPF Neighbors:  3
```

show ip route

```
=========================================================================
Ip Route
=========================================================================
DST           MASK            NEXT          COST VLAN PORT PROT TYPE PRF

-------------------------------------------------------------------------
172.3.3.0     255.255.255.252 10.1.1.18       20 100  T#5  O    IB   20
172.2.2.0     255.255.255.0   10.1.1.10       10 101  T#1  O    IB  120
172.1.1.0     255.255.255.0   10.1.1.10       30 101  T#1  O    IB   20
203.203.100.0 255.255.255.0   203.203.100.53   1   1  ---- C    DB    0
20.1.1.0      255.255.255.0   10.1.1.10       10 101  T#1  O    IB  120
10.1.1.24     255.255.255.252 10.1.1.10       20 101  T#1  O    IB   20
10.1.1.20     255.255.255.252 10.1.1.22        1 102  ---- C    DB    0
10.1.1.16     255.255.255.252 10.1.1.17        1 100  ---- C    DB    0
10.1.1.8      255.255.255.252 10.1.1.9         1 101  ---- C    DB    0
10.1.1.0      255.255.255.252 10.1.1.10       30 101  T#1  O    IB   20
Total Routes: 10
-------------------------------------------------------------------------
TYPE Legend: I=Indirect Route, D=Direct Route, A=Alternative Route,
B=Best Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R3 Status

show vlan

```
Id   Name       Type Protocol User PID Active IVL/SVL Mgmt
---  ---------  ---- -------- -------- ------ ------- ----
1    VLAN #1    Port None     0x0000   Yes    IVL     Yes
Port Members: 4-6,9,12,15-26
103  VLAN #103  Port None     0x0000   Yes      IVL     No
Port Members: 7-8
104  VLAN #104  Port None     0x0000   Yes    IVL     No
Port Members: 1-2
105  VLAN #105  Port None     0x0000   Yes    IVL     No
Port Members: 13-14
1001 VLAN #1001 Port None     0x0000   Yes    IVL     No
Port Members: 10
```

show vlan ip

```
Id  ifIndex Address         Mask             MacAddress    Offset Routing
1   10001   203.203.100.52 255.255.255.0   00:15:9B:F1:FC:40 1   Enabled
103 10103   10.1.1.2       255.255.255.252 00:15:9B:F1:FC:42 3   Enabled
104 10104   10.1.1.25      255.255.255.252 00:15:9B:F1:FC:43 4   Enabled
105 10105   20.1.1.1       255.255.255.0   00:15:9B:F1:FC:44 5   Enabled
1001 11001  172.1.1.1      255.255.255.0   00:15:9B:F1:FC:41 2   Enabled
```

show ip rip

```
Default Import Metric:  8
Domain:
HoldDown Time:  120
Queries:  0
Rip:  Enabled
Route Changes:  1
Timeout Interval:  180
Update Time:  30
```

show ip rip interface

```
IP Address      Enable Send           Receive      Advertise When Down
--------------- ------ -------------- ------------ -------------------
10.1.1.2        false  rip1Compatible rip1OrRip2   false
10.1.1.25       false  rip1Compatible rip1OrRip2   false
20.1.1.1        true   rip1Compatible rip1OrRip2   false
172.1.1.1       false  rip1Compatible rip1OrRip2   false
203.203.100.52  false  rip1Compatible rip1OrRip2   false

RIP   Dflt   Dflt   Trigger AutoAgg
IP Address   Cost Supply Listen Update  Enable Supply Listen PoisonProxy
----------   ---- ------ ------ ------- ------ ------ ------ -----------
10.1.1.2       1 false  false  false   false   true   true   false false
10.1.1.25      1 false  false  false   false   true   true   false false
20.1.1.1       1 false  false  false   false   true   true   false false
172.1.1.1      1 false  false  false   false   true   true   false false
203.203.100.52 1 false  false  false   false   true   true   false false

IP Address      RIP In Policy
--------------  -------------
10.1.1.2
10.1.1.25
20.1.1.1
172.1.1.1
203.203.100.52

IP Address      RIP Out Policy
--------------  --------------
10.1.1.2
10.1.1.25
20.1.1.1        Allow
172.1.1.1
203.203.100.52

IP Address      Holddown Timeout
--------------- -------- -------
10.1.1.2        120      180
10.1.1.25       120      180
20.1.1.1        120      180
172.1.1.1       120      180
203.203.100.52  120      180
```

show route-map detail

```
=================================================================
Route Policy
=================================================================
Name  Allow,   Id 1,   Seq 1
-----------------------------------------------------------------
```

show route-map detail

```
Match:
enable : enable
mode : permit
match-protocol : direct,ospf
match-interface :
match-metric : 0
match-network :
match-next-hop :
match-route-type : any
match-route-src :
Set:
set-injectlist :
set-mask : 0.0.0.0
set-metric : 5
set-metric-type : type2
set-nssa-pbit : enable
set-metric-type-internal : 0
set-preference : 0
--------
Router ID:  1.1.1.3
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  False
AS Boundary Router Config Status:  True
External Link-State Advertisements:  2
External Link-State Checksum:  49786(0xc27a)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  9
New Link-State Advertisements Received:  39
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
---------------------------------------------------------
```

show ip ospf redistribute

```
Source Metric Metric Type Subnet   Enabled Route Policy
------ ------ ----------- -------- ------- ------------
Direct 10     Type 2      Allow    True
RIP    10     Type 2      Allow    True
```

show ip ospf

```
Router ID: 1.1.1.3
Admin Status: Enabled
Version Number: 2
Area Border Router Oper Status: False
AS Boundary Router Config Status: True
External Link-State Advertisements: 2
External Link-State Checksum: 49786(0xc27a)
Type-of-Service (TOS) Routing Supported: False
Originated Link-State Advertisements: 9
New Link-State Advertisements Received: 39
OSPF Traps: Disabled
Auto Virtual Link Creation: Disabled
SPF Hold-Down Time: 10
RFC 1583 Compatibility: Enabled
```

show ip ospf area

```
Area ID:  0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  1
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID:  0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  4
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  448840(0x6d948)
```

show ip ospf

```
Interface:  10.1.1.2
Area ID:  0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  10.1.1.1
Backup Designated Router:  10.1.1.2
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.25
Area ID:  0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  10.1.1.26
Backup Designated Router:  10.1.1.25
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  20.1.1.1
Area ID:  0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  172.1.1.1
Area ID:  0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  172.1.1.1
Backup Designated Router:  0.0.0.0
Authentication Type:  None
```

show ip ospf

```
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  203.203.100.52
Area ID:  0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf neighbor

```
Interface  Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
---------  ------------- -------------- --- ----- ----------- ----
10.1.1.2   1.1.1.1       10.1.1.1       100 Full  0           Dyn
10.1.1.25  1.1.1.4       10.1.1.26      1   Full  0           Dyn
Total OSPF Neighbors:  2
```

show ip route

```
========================================================================
Ip Route
========================================================================
DST           MASK            NEXT          COST VLAN PORT PROT TYPE PRF
------------------------------------------------------------------------
172.2.2.0     255.255.255.0   20.1.1.2      2    105  T#4  R    IB   100
172.3.3.0     255.255.255.252 10.1.1.1      40   103  T#1  O    IB    25
172.1.1.0     255.255.255.0   172.1.1.1     1    1001 ---- C    DB     0
20.1.1.0      255.255.255.0   20.1.1.1      1    105  ---- C    DB     0
10.1.1.16     255.255.255.252 10.1.1.1      30   103  T#1  O    IB    25
10.1.1.20     255.255.255.252 10.1.1.1      20   103  T#1  O    IB    25
10.1.1.24     255.255.255.252 10.1.1.25     1    104  ---- C    DB     0
10.1.1.8      255.255.255.252 10.1.1.26     20   104  T#2  O    IB    20
10.1.1.0      255.255.255.252 10.1.1.2      1    103  ---- C    DB     0
Total Routes: 9
------------------------------------------------------------------------
TYPE Legend: I=Indirect Route, D=Direct Route, A=Alternative Route,
B=Best Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R4 Status

show vlan

```
Id   Name       Type Protocol User PID Active IVL/SVL Mgmt
---  ---------  ---- -------- -------- ------ ------- ----
1    VLAN #1    Port None     0x0000   Yes    IVL     Yes
Port Members: 3-6,9-26
101  VLAN #101  Port None     0x0000   Yes    IVL     No
Port Members: 7-8
```

show vlan

```
104  VLAN #104  Port None      0x0000    Yes     IVL      No
Port Members: 1-2
```

show vlan ip

```
Id  ifIndex Address         Mask            MacAddress        Offset Routing
1   10001   203.203.100.54 255.255.255.0   00:15:9B:F2:2C:40 1      Enabled
101 10101   10.1.1.10       255.255.255.252 00:15:9B:F2:2C:41 2      Enabled
104 10104   10.1.1.26       255.255.255.252 00:15:9B:F2:2C:42 3      Enabled
```

show ip ospf

```
Router ID:  1.1.1.4
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  False
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  45698(0xb282)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  5
New Link-State Advertisements Received:  34
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID:  0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  1
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID:  0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  3
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  409758(0x6409e)
```

show ip ospf interface

```
Interface:  10.1.1.10
Area ID:  0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  10.1.1.9
Backup Designated Router:  10.1.1.10
Authentication Type:  None
MTU Ignore:  Yes
```

show ip ospf interface

```
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.26
Area ID:  0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  10.1.1.25
Backup Designated Router:  10.1.1.26
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  203.203.100.54
Area ID:  0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf neighbor

```
Interface Nbr Router ID Nbr IP Address  Pri State     RetransQLen Perm
--------- ------------- --------------- --- --------- ----------- ----
10.1.1.10 1.1.1.2       10.1.1.9        50  Full      0           Dyn
10.1.1.26 1.1.1.3       10.1.1.25       1   Full      0           Dyn
Total OSPF Neighbors:  2
```

show ip route

```
========================================================================
Ip Route
========================================================================
DST            MASK            NEXT          COST VLAN PORT PROT TYPE PRF
------------------------------------------------------------------------
172.2.2.0      255.255.255.0   10.1.1.25     10   104  T#2  O    IB   120
172.3.3.0      255.255.255.252 10.1.1.9      30   101  T#1  O    IB    25
172.1.1.0      255.255.255.0   10.1.1.25     20   104  T#2  O    IB    20
20.1.1.0       255.255.255.0   10.1.1.25     10   104  T#2  O    IB   120
10.1.1.16      255.255.255.252 10.1.1.9      20   101  T#1  O    IB    25
10.1.1.20      255.255.255.252 10.1.1.9      20   101  T#1  O    IB    25
10.1.1.24      255.255.255.252 10.1.1.26      1   104  ---- C    DB     0
10.1.1.8       255.255.255.252 10.1.1.10      1   101  ---- C    DB     0
10.1.1.0       255.255.255.252 10.1.1.25     20   104  T#2  O    IB    20
Total Routes: 9
------------------------------------------------------------------------
TYPE Legend: I=Indirect Route, D=Direct Route, A=Alternative Route,
B=Best Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R5 Status

show vlan

```
Id   Name       Type Protocol User PID Active IVL/SVL Mgmt
---  ---------- ---- -------- -------- ------ ------- ----
1    VLAN #1    Port None     0x0000   Yes    IVL     Yes
Port Members: 24-26
100  VLAN #100  Port None     0x0000   Yes    IVL     No
Port Members: 5-6
1000 VLAN #1000 Port None     0x0000   Yes    IVL     No
Port Members: 10
```

show vlan ip

```
Id  ifIndex Address        Mask            MacAddress        Offset Routing
1   10001   203.203.100.51 255.255.255.0   00:15:9B:F8:1C:40 1      Enabled
100 10100   10.1.1.18      255.255.255.252 00:15:9B:F8:1C:41 2      Enabled
1000 11000   172.3.3.1     255.255.255.252 00:15:9B:F8:1C:42 3      Enabled
```

show ip ospf

```
Router ID:  1.1.1.5
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  False
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  48
New Link-State Advertisements Received:  387
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID:  0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  3
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID:  0.0.0.2
Import Summaries:  Yes
Import Type:  No
External Intra-Area SPF Runs:  11
Reachable Area Border Routers:  1
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  9
Link-State Advertisements Checksum:  274851(0x431a3)
Stub Metric:  1
Stub Metric Type:  OSPF Metric
```

show ip ospf interface

```
Interface:  10.1.1.18
Area ID:  0.0.0.2
Admin State:  Enabled
Type:  Broadcast
Priority:  0
Designated Router:  10.1.1.17
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  172.3.3.1
Area ID:  0.0.0.2
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  172.3.3.1
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  203.203.100.51
Area ID:  0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf

```
Interface Nbr Router ID   Nbr IP Address   Pri State     RetransQLen Perm
--------- --------------- ---------------- --- --------- ----------- ----
10.1.1.18 1.1.1.2         10.1.1.17        50  Full      0           Dyn
Total OSPF Neighbors:  1
```

show ip route

```
========================================================================
Ip Route
========================================================================
DST            MASK            NEXT          COST VLAN PORT PROT TYPE PRF
------------------------------------------------------------------------
172.3.3.0      255.255.255.252 172.3.3.1      1   1000 ---- C    DB     0
172.1.1.0      255.255.255.0   10.1.1.17     40   100  T#5  O    IB    25
10.1.1.16      255.255.255.252 10.1.1.18      1   100  ---- C    DB     0
10.1.1.24      255.255.255.252 10.1.1.17     30   100  T#5  O    IB    25
10.1.1.20      255.255.255.252 10.1.1.17     20   100  T#5  O    IB    25
10.1.1.8       255.255.255.252 10.1.1.17     20   100  T#5  O    IB    25
10.1.1.0       255.255.255.252 10.1.1.17     40   100  T#5  O    IB    25
0.0.0.0        0.0.0.0         10.1.1.17     11   100  T#5  O    IB    25
Total Routes: 8
------------------------------------------------------------------------
```

show ip route
`TYPE Legend: I=Indirect Route, D=Direct Route, A=Alternative Route,` `B=Best Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW`

Diagnosing neighbor state problems

At initial startup, routers transmit hello packets in an attempt to find other OSPF routers with which form adjacencies. After the hello packets are received, the routers perform an initialization process, which causes the routers to transition through various states before the adjacency is established. The following table lists the states a router can go through during the process of forming an adjacency.

Table 10: OSPF neighbor states

Step	State	Description
1	Down	Indicates that a neighbor was configured manually, but the router did not received any information from the other router. This state can occur only on NBMA interfaces.
2	Attempt	On an NBMA interface, this state occurs when the router attempts to send unicast hellos to any configured interfaces. The Avaya Ethernet Routing Switch 4000 Series does not support NBMA type.
3	Init	The router received a general hello packet (without its Router ID) from another router.
4	2-Way	The router received a Hello directed to it from another router. (The hello contains its Router ID)
5	ExStart	Indicates the start of the Master/Slave election process.
6	Exchange	Indicates the link state database (LSDB) is exchanged
7	Loading	Indicates the processing state of the LSDB for input into the routing table. The router can request LSA for missing or corrupt routes.
8	Full	Indicates the normal full adjacency state.

OSPF neighbor state information

Neighbor state information can be accessed by using the **show ip ospf neighbor** command.

```
ERS4000#show ip ospf neighbor
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------- -------------- --- ----- ----------- ----
10.1.1.22 1.1.1.1       10.1.1.21      100 Full  0           Dyn
10.1.1.17 1.1.1.5       10.1.1.18      0   Full  0           Dyn
10.1.1.9  1.1.1.4       10.1.1.10      1   Full  0           Dyn
```

Problems with OSPF occur most often during the initial startup, when the router cannot form adjacencies with other routers and the state is stuck in the **Init** or **ExStart/Exchange** state.

Init State Problems

A router can become stuck in an **Init** state and not form adjacencies. There are several possible causes for this problem:

- Authentication mismatch or configuration problem
- Area mismatch for Stub or NSSA
- Area ID mismatch
- Hello Interval or Dead Interval mismatch

To determine any mismatches in OSPF configuration, use the **show ip ospf ifstats mismatch** command.

ExStart/Exchange problems

Even though routers can recognize each other and have moved beyond two way communications, routers can become stuck in the **ExStart/Exchange** state.

A mismatch in maximum transmission unit (MTU) sizes between the routers usually causes this type of problem. For example, one router could be set for a high MTU size and the other router a smaller value. Depending on the size of the link state database, the router with the smaller value may not be able to process the larger packets and thus be stuck in this state. To avoid this problem, ensure that the MTU size value for both routers match. This problem is usually encountered during interoperations in networks with other vendor devices.

Note:

The Avaya Ethernet Routing Switch 4000 Series automatically checks for OSPF MTU mismatches.

In the Avaya Ethernet Routing Switch 4000 Series , the supported MTU size for OSPF is 1500 bytes by default. Incoming OSPF database description (DBD) packets are dropped if their MTU size is greater than this value.

Chapter 9: RIP configuration using ACLI

This section describes how to configure RIP using ACLI.

RIP is a distance vector protocol used to dynamically discover network routes based on information passed between routers in the network.

Prerequisites

- Enable IP routing globally.
- Assign an IP address to the VLAN or port for which you want to enable RIP.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

Enabling RIP globally

Use this procedure to globally enable RIP on the switch.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Enable RIP on the switch by using the following command:

```
[default] [no] router rip enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Globally disables RIP on the switch.
no	Globally disables RIP on the switch.

Configuring global RIP timers

Use this procedure to set the RIP global timeout, holddown timer, and update timer.

Prerequisites

- Log on to the RIP Router Configuration mode in ACLI.

Procedure steps

Configure the global RIP timers, enter the following from the RIP Router Configuration command mode:

```
[default] timers basic holddown <holdown-timer> timeout
<global-timeout> update <update-timer>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Returns the parameters to the factory default timer values: • holddown timer: 120 seconds • global timeout: 180 seconds • update timer: 30 seconds
<holdown-timer>	Specifies the global holddown timer, which is the length of time (in seconds) that RIP maintains a route in the garbage list after determining that it is unreachable. During this period, RIP continues to advertise the garbage route with a metric of infinity (16). If a valid update for a garbage route is received within the holddown period, the router adds the route back into its routing

Variable	Value
	table. If no update is received, the router deletes the garbage list entry. Range is 0–360 seconds. Default is 120 seconds.
`<global-timeout>`	Specifies the global timeout interval parameter. If a RIP router does not receive an update from another RIP router within the configured timeout period, it moves the routes advertised by the nonupdating router to the garbage list. The timeout interval must be greater than the update timer. Range is 15–259200 seconds. Default is 180 seconds.
`<update-timer>`	Specifies a value for the RIP update timer, which is the time interval (in seconds) between regular RIP updates. The update timer value must be less than the timeout interval. Range is 0–360 seconds. Default is 30 seconds.

Configuring the default RIP metric value

Use this procedure to configure a default metric to apply to routes not learned through RIP but imported into the RIP domain. The switch applies this default metric to redistributed routes if the associated route policy does not specify a metric for the redistributed protocol, such as OSPF. The range is 0 to 15, and the default is 8.

Prerequisites

- Log on to the RIP Router Configuration mode in ACLI.

Procedure steps

Configure the default RIP metric value by using the following command:

```
[default] default-metric <metric_value>
```

Variable definitions

The following table describes the command variables.

Variable	Value
`<metric_value>`	Specifies a metric value between 0 and 15.

Variable	Value
`default`	Returns the switch to the factory default RIP default import metric value (8).

Displaying global RIP information

Use this procedure to display the global RIP configuration.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display the global RIP configuration by using the following command:

```
show ip rip
```

Job aid

The following table shows the field descriptions for the **`show ip rip`** command.

Field	Description
Default Import Metric	Indicates the value of the default import metric.
Domain	Indicates the value inserted into the Routing Domain field of all RIP packets sent on this device. This value is not configurable.
HoldDown Time	Indicates the value of the holddown timer.
Queries	Indicates the number of responses the router has sent in response to RIP queries from other systems.
Rip	Indicates whether RIP is enabled.
Route Changes	Indicates the number of route changes the RIP process has made to the routing database.
Timeout Interval	Indicates the RIP timeout interval.
Update Time	Indicates the value of the RIP update timer.

Configuring RIP on an interface

Use this procedure to configure RIP parameters on an interface.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure RIP for an interface by using the following command:

```
[default] [no] ip rip [advertise-when-down enable] [auto-aggregation enable] [cost <cost>] [default-listen enable] [default-supply enable] [enable] [holddown <holddown> | <global>] [listen enable] [poison enable] [proxy-announce enable] [receive version {rip1 | rip1orrip2 | rip 2}] [send version {notsend | rip1 | rip1comp | rip 2}] [supply enable] [timeout {<timeout> | global}] [triggered enable]
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Sets the specified parameter to the default value.
no	Removes or disables the specified configuration.
advertise-when-down enable	Enables RIP advertisements for an interface even when the link to the network fails. The router continues to advertise the subnet even if that particular network is no longer connected (no link in the enabled VLAN). This feature does not advertise the route until the VLAN is first enabled. After the VLAN is enabled, the route is advertised even when the link fails. By default, advertise when down functionality is disabled.
auto-aggregation enable	Enables auto aggregation on the RIP interface. After you enable auto aggregation, the Ethernet Routing Switch automatically aggregates routes to their natural net mask when they are advertised on an

Variable	Value
	interface in a network of a different class. Automatic route aggregation can be enabled only in RIP2 mode or RIP1 compatibility mode. By default, auto aggregation is disabled.
`cost <cost>`	Specifies the RIP cost (metric) for this interface in a range from 1 to 15. The default cost is 1.
`default-listen enable`	Enables the interface to accept default routes learned through RIP updates. The default setting is disabled.
`default-supply enable`	Enables the interface to send default route information in RIP updates. This setting takes effect only if a default route exists in the routing table. The default setting is disabled.
`enable`	Enables RIP on the interface.
`holddown <holddown> | <global>`	Specifies the interface holddown timer, which is the length of time (in seconds) that RIP maintains a route in the garbage list after determining that it is unreachable. During this period, RIP continues to advertise the garbage route with a metric of infinity (16). If a valid update for a garbage route is received within the holddown period, the router adds the route back into its routing table. If no update is received, the router deletes the garbage list entry. • holddown—overrides the global parameter and does not change if the global parameter is modified. Range is 0–360 seconds. • global—default global holddown parameter (120 seconds)
`listen enable`	Enables this interface to listen for RIP advertisements. The default value is enabled.
`poison enable`	Specifies whether RIP routes on the interface learned from a neighbor are advertised back to the neighbor. If poison reverse is disabled, split horizon is invoked and IP routes learned from an immediate neighbor are not advertised back to the neighbor. If poison reverse is enabled, the RIP updates sent to a neighbor from which a route is learned are "poisoned" with a metric of 16. The receiving neighbor ignores this route because the metric 16 indicates infinite hops in the network. By default, poison reverse is disabled.
`proxy-announce enable`	Enables proxy announcements on a RIP interface. When proxy announcements are enabled, the source of a route and its next hop are treated as the same when processing received updates. So, instead of the advertising router being used as the source, the next hop is. Proxy announcements are disabled by default.
`receive version {rip1 | rip1orrip2 | rip 2}`	Specifies the RIP version received on this interface. Default is rip1orrip2.
`send version {notsend |`	Specifies the RIP version sent on an interface. Default is rip1compatible.

Variable	Value
rip1 | rip1comp | rip 2}	
supply enable	Enables RIP route advertisements on this interface. The default value is enabled.
timeout *<timeout>* | *<global>*	Specifies the RIP timeout value on this interface. If a RIP interface does not receive an update from another RIP router within the configured timeout period, it moves the routes advertised by the nonupdating router to the garbage list. The timeout interval must be greater than the update timer. • timeout—sets the interface timeout. Value ranges from 15 to 259200 seconds. • global—sets the timeout to the global default (180 seconds). The interface timer setting overrides the global parameter and does not change if the global parameter is changed.
triggered enable	Enables automatic triggered updates on this RIP interface. Default is disabled.

Displaying the global RIP configuration

Use this procedure to display RIP configuration information for the switch.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display global RIP configuration information by using the following command:

```
show ip rip
```

Displaying RIP interface configuration

Use this procedure to display configuration for a RIP interface.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display RIP interface configuration by using the following command:

show ip rip interface [*<vid>*] [Ethernet *<portlist>*] [vlan *<vid>*]

Variable definitions

The following table describes the command variables.

Variable	Value
[*<vid>*]	Displays RIP information for the specified VLAN.
[Ethernet *<portlist>*]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[vlan *<vid>*]	Displays RIP information for VLAN interfaces only. If no VLAN ID is specified, all VLAN information is displayed.

Job aid

The following table shows the field descriptions for the **show ip rip interface** command.

Field	Description
unit/port	Indicates the unit and port of the RIP interface.

Field	Description
IP Address	Indicates the IP address of the RIP interface.
Enable	Indicates whether RIP is enabled or disabled on the interface.
Send	Indicates which send mode is enabled.
Receive	Indicates which receive mode is enabled.
Advertise When Down	Indicates whether the advertise when down feature is enabled.
RIP Cost	Indicates the RIP cost (metric) for this interface.
Dflt Supply	Indicates whether the interface sends the default route in RIP updates, if a default route exists in the routing table.
Dflt Listen	Indicates whether the interface listens for default routes in RIP updates.
Trigger Update	Indicates whether triggered updates are enabled.
AutoAgg Enable	Indicates whether auto aggregation is enabled.
Supply	Indicates whether the interface is enabled to supply updates for RIP.
Listen	Indicates whether the interface is enabled to listen for RIP routes.
Poison	Indicates whether RIP routes on the interface learned from a neighbor are advertised back to the neighbor.
Proxy	Indicates whether proxy announcements are enabled.
RIP IN Policy	Indicates the RIP policy for inbound filtering on the interface.
RIP Out Policy	Indicates the RIP policy for outbound filtering on the interface.
Holddown	Indicates the value of the RIP holddown timer for the interface.
Timeout	Indicate the RIP timeout interval for the interface.

Manually triggering a RIP update

Use this procedure to manually trigger a RIP update on an interface.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

Manually trigger a RIP update by using the following command:

```
manualtrigger ip rip interface vlan <vid>
```

Chapter 10: RIP configuration examples using ACLI

This section provides examples to help you create common RIP configurations.

You can configure RIP on a VLAN or brouter port basis.

Note:

In many of the following configuration examples, a brouter port is used to create a connection to the network core. You can also use L3 enabled VLAN interfaces instead of brouter ports to create these connections.

RIP configuration tasks

To perform a basic RIP configuration on a VLAN, perform the following steps.

1. Configure the interface, assign an IP address and add ports.

```
ERS4000#enable
ERS4000#config terminal
ERS4000(config)#vlan create 51 name "VLAN-51" type port
ERS4000(config)#interface vlan 51
ERS4000(config-if)#ip address 10.10.1.1 255.255.255.0
ERS4000(config-if)#exit
ERS4000(config)#vlan members add 51 8-9
```

2. Enable RIP using one of the following command sequences.

```
ERS4000(config)#interface vlan 51
ERS4000(config-if)#ip rip enable
ERS4000(config-if)#exit
```

OR

```
ERS4000(config)#router rip
ERS4000(config-router)#network 10.10.1.1
ERS4000(config-router)#exit
```

3. Select the VLAN to configure RIP interface properties.

```
ERS4000(config)#interface vlan 51
```

4. Disable Supply RIP Updates on the VLAN, if required.

```
ERS4000(config-if)#no ip rip supply enable
```

5. Disable Listen for RIP Updates on the VLAN, if required.

```
ERS4000(config-if)#no ip rip listen enable
```

6. Enable Default Route Supply on the VLAN, if a default route exists in the route table.

   ```
   ERS4000(config-if)#ip rip default-supply enable
   ```

7. Enable Default Route Listen on the VLAN to add a default route to the route table, if advertised from another router.

   ```
   ERS4000(config-if)#ip rip default-listen enable
   ```

8. Add the Out Route Policy to the VLAN (this step assumes that you have previously configured the route policy).

   ```
   ERS4000(config-if)#ip rip out-policy map1
   ```

9. Enable Triggered Updates on the VLAN, if required.

   ```
   ERS4000(config-if)#ip rip triggered enable
   ```

10. Configure the cost of the VLAN link by entering a value of 1 to 15; where 1 is the default.

    ```
    ERS4000(config-if)#ip rip cost 2
    ```

11. Configure send mode parameters on the VLAN.

    ```
    ERS4000(config-if)#ip rip send version rip2
    ```

12. Configure receive mode parameters on the VLAN.

    ```
    ERS4000(config-if)#ip rip receive version rip2
    ```

13. Enable poison reverse on the VLAN.

    ```
    ERS4000(config-if)#ip rip poison enable
    ```

Configuring RIP

This section describes the set up of a basic RIP configuration between two Avaya Ethernet Routing Switch 4000 Series routers. As shown in the following diagram, router ERS2 is configured between router ERS1 and the edge of the network core. Two VLANs (VLAN 2 and 3) are associated with ERS1.

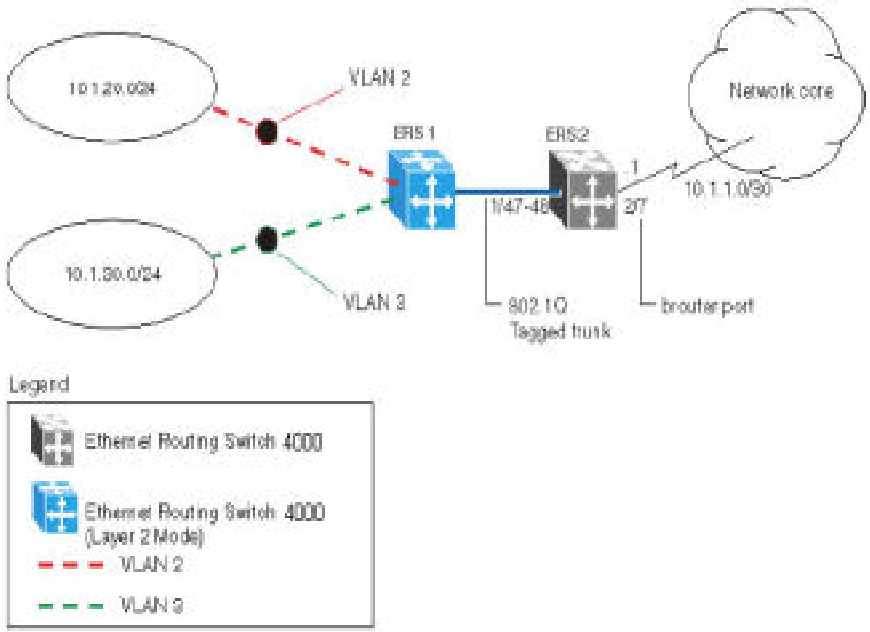

Figure 34: RIP configuration example

In this example:

- ERS1 is an edge switch with two configured VLANs, VLAN 2 and 3. It is connected to aggregation switch ERS2 on ports 1/47 and 1/48.
- Port 2/7 of ERS2 is configured as a RIP enabled brouter port to connect to the network core.

Use the following procedure to configure router RIP as illustrated in the preceding drawing:

1. Configure tagging on ports 1/47 and 1/48.

   Tagging is required to support multiple VLANs on the same interface.

   Example

   ```
   ERS4000#enable
   ERS4000#config terminal
   ERS4000(config)#vlan ports 1/47-48 tagging tagAll
   ```

2. Configure ERS2 for VLAN 2 access.

   Create a port-based VLAN (VLAN 2) using spanning tree group 1 and include ports 1/47 and 1/48 in VLAN 2.

   Example

   ```
   ERS4000(config)#vlan create 2 name "VLAN-2" type port
   ERS4000(config)#vlan member add 2 port 1/47-48
   ```

3. Assign the IP address 10.1.20.2/24 to VLAN 2.

Example

```
ERS4000(config)#interface vlan 2
ERS4000(config-if)#ip address 10.1.20.2 255.255.255.0
```

4. Enable RIP for VLAN 2 and disable RIP supply and listen. RIP supply and listen are not required because no router is attached to VLAN 2.

   Example

```
ERS4000(config)#interface vlan 2
ERS4000(config-if)#ip rip enable
ERS4000(config-if)#ip rip supply disable
ERS4000(config-if)#ip rip listen disable
```

5. Configure ERS2 for VLAN 3 access

   Create a port-based VLAN (VLAN 3) using spanning tree group 1 and include ports 1/47 and 1/48 in VLAN 3.

   Example

```
ERS4000(config)#vlan create 3 name "VLAN-3" type port
ERS4000(config)#vlan member add 3 port 1/47-48
```

6. Assign the IP address 10.1.30.2/24 to VLAN 3.

   Example

```
ERS4000(config)#interface vlan 3
ERS4000(config-if)#ip address 10.1.30.2 255.255.255.0
```

7. Enable RIP for VLAN 3 and disable RIP supply and listen. RIP supply and listen are not required because no router is attached to VLAN 3.

   Example

```
ERS4000(config)#interface vlan 3
ERS4000(config-if)#ip rip enable
ERS4000(config-if)#ip rip supply disable
ERS4000(config-if)#ip rip listen disable
```

8. Configure brouter port 2/7 on ERS2.

   a. Assign the IP address 10.1.1.1/30 to port 2/7 using brouter VLAN 2090.

      Example

```
ERS4000(config)# interface Ethernet 2/7
ERS4000(config-if)# brouter vlan 2090 subnet 10.1.1.1/30
```

      Note:

      Use of the brouter command above requires the use of Variable Length Subnetting. Use of a dotted decimal subnet mask is not allowed.

   b. Enable RIP on the interface.

Example

```
ERS4000(config)# interface Ethernet 2/7
ERS4000(config-if)# ip rip enable
```

9. Enable IP routing and RIP globally.

   Example

```
ERS4000(config)#ip routing
ERS4000(config)#router rip enable
```

A list of the commands used to create this configuration can be displayed using the **show running-config** command. Using this command on ERS2 would list the following commands:

```
! *** VLAN *** !
vlan igmp unknown-mcast-no-flood disable
vlan configcontrol strict
auto-pvid
vlan name 1 "VLAN #1"
vlan create 2 name "VLAN-2" type port
vlan create 3 name "VLAN-3" type port
vlan members 2 1/47-48
vlan members 3 1/47-48
! *** RIP ***  !
router rip
router rip enable
timers basic holddown 120
timers basic timeout 180 update 30 default-metric 8
network 10.1.20.2
network 10.1.30.2
network 10.1.1.1
interface vlan 2
no ip rip listen enable
no ip rip supply enable
interface vlan 3
no ip rip listen enable
no ip rip supply enable
! *** Brouter Port *** !
interface Ethernet ALL
brouter port 2/7 vlan 3 subnet 10.1.1.1/30
```

The following commands can be used to confirm the configuration of RIP parameters:

Command	Description
show vlan	This command is used to display information about the currently configured switch VLANs.
show vlan ip	This command is used to display IP address information about VLANs that have been assigned addresses on the switch.
show ip rip	This command displays information on the global switch RIP configuration.
show ip route	This command displays the switch routing table.
show ip rip interface	This command displays information about the RIP interfaces present on the switch.

Configuring RIP version 2

When RIP is enabled on an interface, it operates by default in **rip1compatible** send mode and **rip1orRip2** receive mode. Depending on configuration requirements, the Avaya Ethernet Routing Switch 4000 Series can be configured to operate using RIP version 1 or 2. The configuration illustrated below demonstrates an Avaya Ethernet Routing Switch 4000 Series switch that has been configured to operate use RIP version 2 only.

This example builds on the previous RIP configuration.

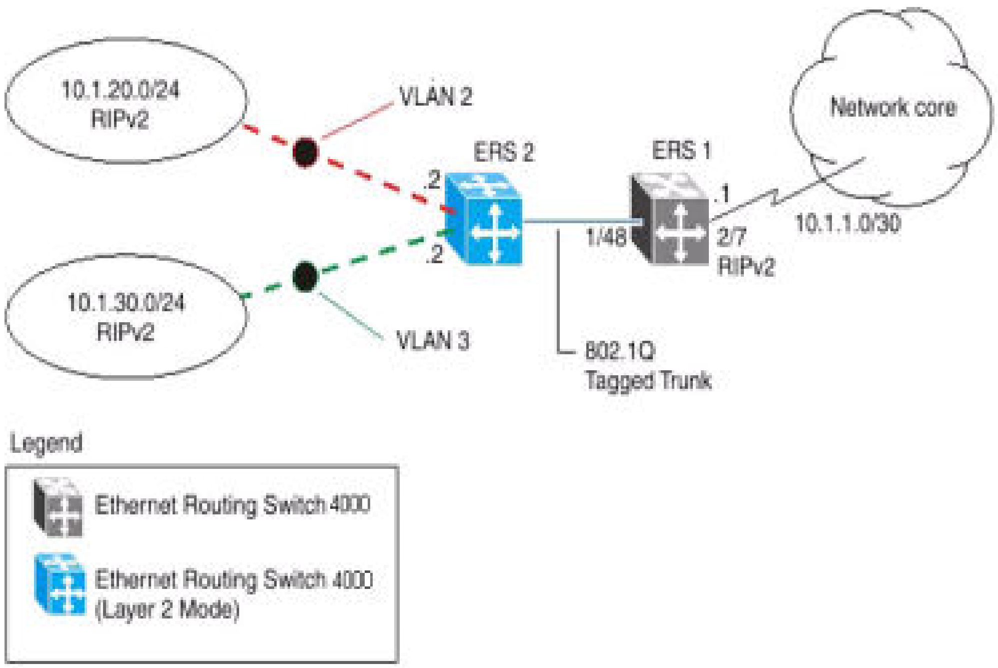

Figure 35: RIPv2 configuration example

Use the following procedure to configure ERS2 to add RIP version 2 to VLAN 2, VLAN 3, and the brouter port..

1. Configure RIP version 2 on VLAN 2. Enable RIP version 2 mode on the IP address used for VLAN 2.

   Example

   ```
   ERS4000#enable
   ERS4000#config terminal
   ERS4000(config)#router rip enable
   ERS4000(config)#interface vlan 2
   ```

```
ERS4000(config-if)#ip rip send version rip2
ERS4000(config-if)#ip rip receive version rip2
```

2. Configure RIP version 2 on VLAN 3. Enable RIP version 2 mode on the IP address used for VLAN 3.

   Example

```
ERS4000(config)#router rip enable
ERS4000(config)#interface vlan 3
ERS4000(config-if)#ip rip send version rip2
ERS4000(config-if)#ip rip receive version rip2
```

3. Configure RIP version 2 on the brouter port. Enable RIP version 2 mode on the IP address used for the brouter port.

   Example

```
ERS4000(config)#router rip enable
ERS4000(config)# interface Ethernet 2/7
ERS4000(config-if)# ip rip enable
ERS4000(config-if)# ip rip send version rip2
ERS4000(config-if)# ip rip receive version rip2
```

Using RIP accept policies

RIP accept policies are used on the Avaya Ethernet Routing Switch 4000 Series to selectively accept routes from RIP updates. If no policies are defined, the default behavior is applied. This default behavior is to add all learned routes to the route table. RIP accept policies are used to:

- Listen to RIP updates only from certain gateways.
- Listen only for specific networks.
- Assign a specific mask to be included with a network in the routing table (such as a network summary).

In the configuration illustrated below, the Avaya Ethernet Routing Switch 4000 Series (ERS1) is configured with a RIP accept policy. This creates a single route directed to ERS3 for all networks configured on it. The accept policy accepts any network from 10.1.240.0 to 10.1.255.0, and creates a single entry in the routing table on ERS1.

A summary route is calculated by comparing the common bits in the address range to derive the summary address. For example, if the range of IP addresses is from 10.1.240.0 to 10.1.255.0:

1. Determine the third octet of the first address: 10.1.240.0 = 1111 0000.
2. Determine the third octet of the ending address: 10.1.255.0 = 1111 1111.
3. Extract the common bits: 240 = 1111 0000 255 = 1111 1111 1111 = 20 bit mask.

Therefore, the network address to use for this example is 10.1.240.0/20

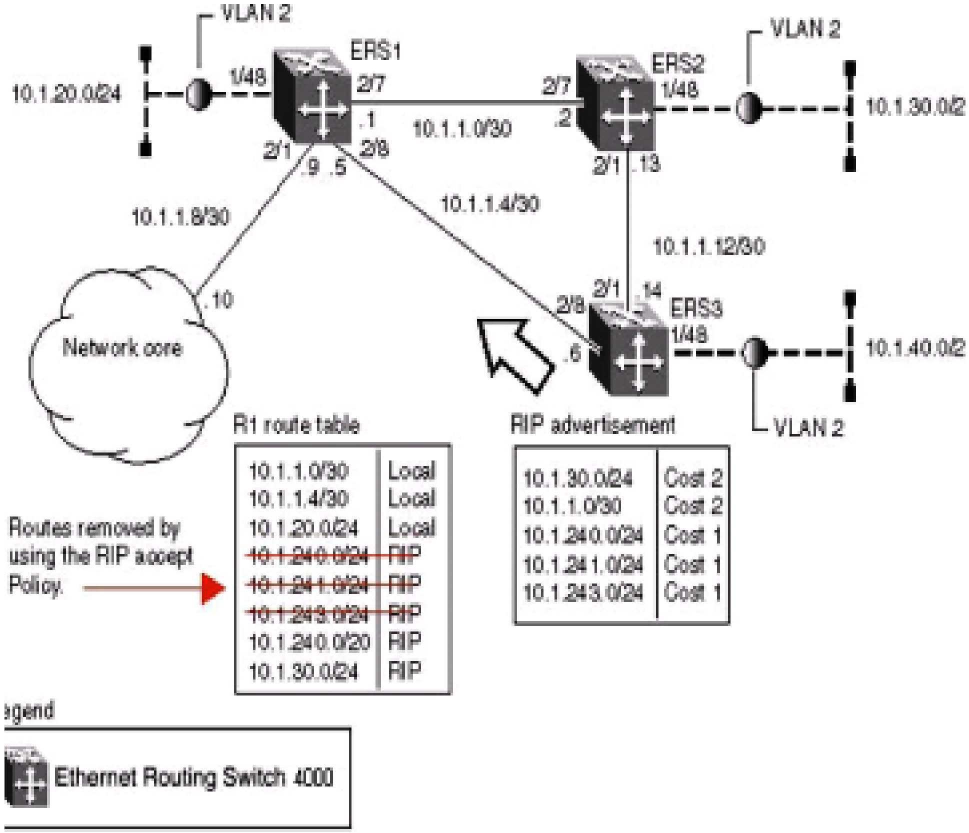

Figure 36: Accept policy configuration

Use the following steps to recreate the above configuration example:

1. Configure the IP prefix list on ERS1.

   Create a prefix list named **Prefix_1** with an IP range from 10.1.240.0 to 10.1.255.0.

   ```
   ERS4000(config)# ip prefix-list Prefix_1 10.1.240.0/20 ge 20
   le 32
   ```

2. Configure the route policy named **rip_pol_1** with match criteria using the IP prefix configured in step 1. This injects one route of 10.1.240.0/20 into the route table.

   ```
   ERS4000(config)# route-map rip_pol_1 1
   ERS4000(config)# route-map rip_pol_1 1 enable
   ERS4000(config)# route-map rip_pol_1 permit 1 enable
   ERS4000(config)# route-map rip_pol_1 permit 1 match network Prefix_1
   ERS4000(config)# route-map rip_pol_1 permit 1 set injectlist Prefix_1
   ```

3. Create brouter port, enable RIP and add route policy to brouter port.

   ```
   ERS4000(config)#interface Ethernet 2/8
   ERS4000(config-if)#brouter port 2/8 vlan 2091 subnet 10.1.1.5/30
   ERS4000(config-if)#ip rip enable
   ```

```
ERS4000(config-if)#ip rip in-policy rip_pol_1
ERS4000(config-if)#exit
```

The **show running-config** command is used to display the current configuration of a switch. Using this command on the above configuration would yield the following results:

Example

```
! *** VLAN ***
! vlan 2091 is brouter
vlan configcontrol flexible
vlan members 1 1-5,7-48
vlan configcontrol automatic
! *** Brouter Port ***
interface Ethernet ALL
brouter port 2/8
vlan 2091 subnet 10.1.1.5/30
exit
! --- Route Policies ---
ip prefix-list Prefix_1 10.1.240.0/20 le 32
route-map rip_pol_1 1
route-map rip_pol_1 1 enable
route-map rip_pol_1 1 set injectlist Prefix_1
! --- RIP ---
interface vlan 2091
ip rip in-policy rip_pol_1
ip rip enable
exit
```

Using RIP announce policies

In the previous configuration example, a RIP accept policy is used on ERS1 to insert a single route into its route table for all networks from ERS3. Instead of using an accept policy on ERS1, a RIP announce policy on ERS3 could be used to announce a single route to both ERS1 and ERS2 for the local network range.

To configure the RIP announce policy on ERS3, use the following configuration steps:

1. Configure the IP prefix list on ERS3 named **Prefix_1** with the IP address 10.1.240.0.

   ```
   ERS4000(config)# ip prefix-list Prefix_1 10.1.240.0/20 ge 20
   le 32
   ```

2. Configure the route policy named **Policy_Rip** with match criteria using the IP prefix configured in step 1.

   ```
   ERS4000(config)# route-map rip_pol_1 1
   ERS4000(config)# route-map rip_pol_1 1 enable
   ```

```
ERS4000(config)# route-map rip_pol_1 permit 1 enable
ERS4000(config)# route-map rip_pol_1 permit 1 set-injectlist Prefix_1
```

3. Add the route policy created in step 2 to VLAN 4.

```
ERS4000(config)#interface vlan 4
ERS4000(config-if)#ip address 10.1.1.1/30
ERS4000(config-if)#ip rip enable
ERS4000(config-if)#ip rip out-policy rip_pol_1
```

To limit the advertising of routes using the announce policy from the routing table, a route policy should be created to deny the route. To configure the RIP announce policy with a limited announce policy on ERS3, use the following configuration steps:

1. Configure the IP prefix list named **`Prefix_2`** with the IP address 10.1.240.0.

```
ERS4000(config)# ip prefix-list Prefix_2 10.1.240.0/20 ge 20
le 20
```

2. Configure the IP route policy named **`rip_pol_2`** with match criteria using the IP prefix configured in Step 1.

```
ERS4000(config)# route-map rip_pol_2 deny 1 enable match network Prefix_2
ERS4000(config)# route-map rip_pol_2 1 match network Prefix_2
```

3. Add the route policy created in step 2 to VLAN 4.

```
ERS4000(config)#interface vlan 4
ERS4000(config-if)#ip address 10.1.1.1/30
ERS4000(config-if)#ip rip enable
ERS4000(config-if)#ip rip out-policy rip_pol_2
```

Chapter 11: VRRP configuration using ACLI

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure that can occur when the single static default gateway router for an end station is lost. This section describes the procedures you can use to configure VRRP on a VLAN using ACLI.

VRRP prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN that you want to enable with VRRP.

Routing is automatically enabled on the VLAN when you assign an IP address to it.

VRRP configuration procedures

To enable VRRP on a VLAN, perform the following steps:

1. Enable VRRP globally on the switch.
2. Assign a virtual router IP address to a virtual router ID.
3. Configure the priority for this router as required.
4. Enable the virtual router.

Configuring global VRRP status using ACLI

Use this procedure to configure the global VRRP status on the switch.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure the VRRP status:

```
[no] router vrrp enable
```

Variable definitions

Variable	Value
[no]	Globally disable VRRP on the switch.

Assigning an IP address to a virtual router ID using ACLI

Use this procedure to associate an IP address with a virtual router ID.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Assign an IP address to a virtual router ID:

```
[no] ip vrrp address <vr_id> <ip_address>
```

Variable definitions

Variable	Value
<ip_address>	The IP address to associate with the virtual router ID
[no]	Removes the IP address from the virtual router ID.
<vr_id>	Specify the virtual router to configure. Value between 1 and 255.

Assigning the router priority for a virtual router ID using ACLI

Use this procedure to assign a priority to the router for specific virtual router ID.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Assign a priority to the router for a specific virtual router ID:

```
ip vrrp <vr_id> priority <priority_value>
```

Variable definitions

Variable	Value
<priority_value>	Specify the priority value for the virtual router ID. Value between 1 and 255.
<vr_id>	Specify the virtual router ID to configure router priority. **Note:** The priority value of 255 is reserved exclusively for IP owners.

Configuring the status of the virtual router using ACLI

Use this procedure to configure the virtual router interface.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Enable or disable the virtual router interface:

```
[no] ip vrrp <vr_id> enable
```

Variable definitions

Variable	Value
[no]	Disables the virtual router.
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP critical IP address using ACLI

Use this procedure to configure the VRRP critical IP address.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure the VRRP critical IP address:

```
[no] ip vrrp <vr_id> critical-ip-addr <ip_address>
```

Variable definitions

Variable	Value
<ip_address>	Specify the critical IP address.
[no]	Removes the configured critical IP address.
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP critical IP status using ACLI

Use this procedure to configure the VRRP critical IP status.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Enable or disable the VRRP critical IP:

```
[no] ip vrrp <vr_id> critical-ip enable
```

Variable definitions

Variable	Value
[no]	Disable the VRRP critical IP.
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP holddown timer using ACLI

Use this procedure to configure the VRRP holddown timer.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure the VRRP holddown timer:

```
ip vrrp <vr_id> holddown-timer <timer_value>
```

Variable definitions

Variable	Value
<timer_value>	Specify the holddown timer value. Value in seconds between 1 and 21600.
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP holddown action using ACLI

Use this procedure to configure the VRRP holddown action.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure the VRRP holddown action:

```
ip vrrp <vr_id> action {none | preempt}
```

Variable definitions

Variable	Value
{none | preempt}	Specify the holddown action. Enter **none** for no action, Enter **preempt** to cancel the holddown timer.
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP advertisement interval using ACLI

Use this procedure to configure the VRRP advertisement interval.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure the VRRP advertisement interval:

```
ip vrrp <vr_id> adver-int <interval>
```

Variable definitions

Variable	Value
<interval>	Specify the advertisement interval in seconds. Value between 1 and 255.
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP fast advertisement interval using ACLI

Use this procedure to configure the VRRP fast advertisement interval.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Configure the VRRP fast advertisement interval:

```
ip vrrp <vr_id> fast-adv-int <interval>
```

Variable definitions

Variable	Value
<interval>	Specify the fast advertisement interval in milliseconds. Value between 200 and 1000.

Variable	Value
<vr_id>	Specify the virtual router ID to configure.

Configuring the VRRP fast advertisement status using ACLI

Use this procedure to enable or disable the VRRP fast advertisement functionality.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Enable or disable VRRP fast advertisement:

```
[no] ip vrrp <vr_id> fast-adv enable
```

Variable definitions

Variable	Value
[no]	Disable the VRRP fast advertisement functionality.
<vr_id>	Specify the virtual router ID to configure.

Configuring ICMP echo replies using ACLI

Use this procedure to enable or disable ICMP echo replies from virtual router IP addresses.

Prerequisites

- Log on to the VRRP router Configuration mode in ACLI.

Procedure steps

Enable or disable ICMP echo replies for VRRP:

```
[no] ping-virtual-address enable
```

Variable definitions

Variable	Value
[no]	Disables ICMP echo replies for VRRP associated addresses.

Displaying VRRP configuration information using ACLI

Use this procedure to display VRRP configuration information. You can display global, address or interface VRRP information.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

1. View global VRRP information:

   ```
   show ip vrrp
   ```

2. View VRRP address information:

   ```
   show ip vrrp address [addr <A.B.C.D>] [vrid <1-255>] [vlan <1-4094>]
   ```

3. View VRRP interface information:

   ```
   show ip vrrp interface [vrid < 1-255>] [vlan <1-4094>] [verbose]
   ```

Variable definitions

Variable	Value
addr <A.B.C.D>	Displays VRRP configuration for the specified IP address.
verbose	Displays additional VRRP configuration information.
vlan <1-4094>	Displays VRRP configuration for the specified VLAN.
vrid <1-255>	Displays VRRP configuration for the specified virtual router ID.

VRRP configuration example 1

The following configuration example shows how to provide VRRP service for two edge host locations.

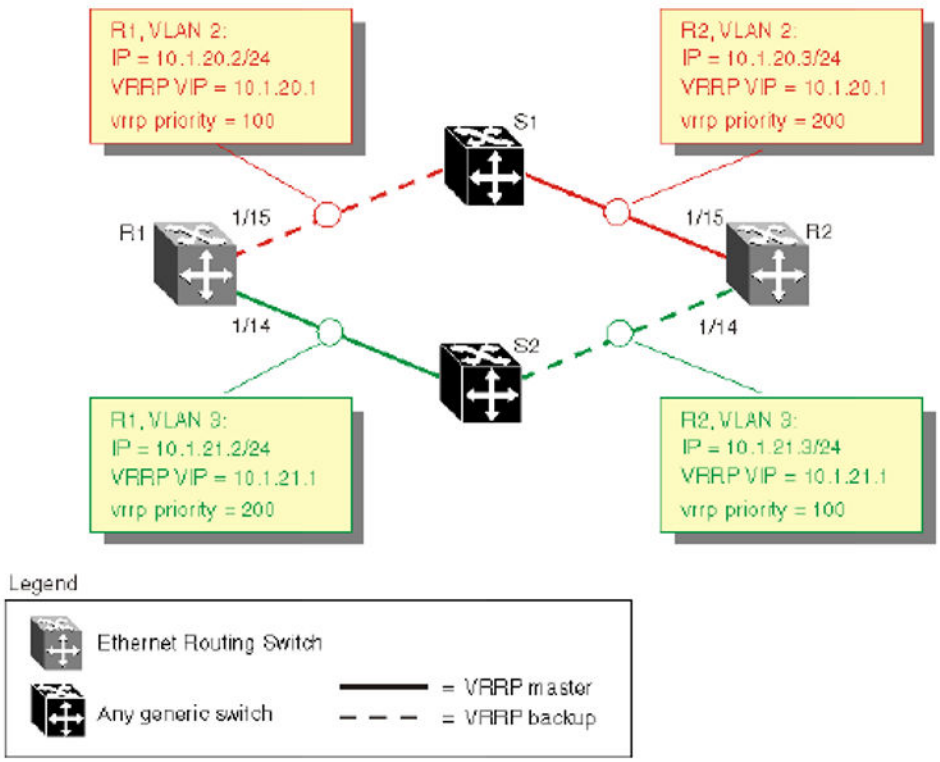

Figure 37: Example VRRP topology 1

In this example, the switches have the following roles:

- R1 is the VRRP master for S2
- R2 is the VRRP master for S1

VRRP is enabled with OSPF as the routing protocol on R1 and R2.

The VRRP priority setting is used to determine which router becomes the VRRP master and which becomes the VRRP backup. In instances where the priority setting is the same for two routers, the higher IP address is the tie breaker. Therefore, it is very important to set the correct VRRP priority. VRRP fast advertisement is enabled in this example to allow for fast failover detection.

The following procedure describes the steps necessary to reproduce the example described above:

1. Configure VLAN 2 on router R1.

    a. Create VLAN 2 on router R1.

    ```
    ERS4000# config terminal
    ERS4000(config)# vlan create 2 type port
    ```

    b. Configure the ports for VLAN 2 on R1.

    ```
    ERS4000# config terminal
    ERS4000(config)# vlan members add 2 1/15
    ```

    c. Configure an IP address for VLAN 2.

    Add IP address 10.1.20.2 / 255.255.255.0 to VLAN 2.

    ```
    ERS4000# config terminal
    ERS4000(config)# interface vlan 2
    ERS4000(config-if)# ip address 10.1.20.2 255.255.255.0
    ```

    d. Configure an OSPF interface for VLAN 2.

    ```
    ERS4000# config terminal
    ERS4000(config)# router ospf enable
    ERS4000(config)# router ospf
    ERS4000(config-router)# network 10.1.20.2
    ```

    e. Configure VRRP on VLAN 2.

    The VRRP VIP address of 10.1.20.1 is added to VLAN 2 using a VRID of 1.

    ```
    ERS4000# config terminal
    ERS4000(config)# router vrrp ena
    ERS4000(config)# interface vlan 2
    ERS4000(config-if)# ip vrrp address 1 10.1.20.1
    ERS4000(config-if)# ip vrrp 1 enable
    ```

Note:

The VRRP priority is not configured here; the priority remains the factory default of 100. Instead, the priority setting on router R2 is set to a higher value when R2 is configured.

Note:

Fast advertisement is disabled by default. Fast advertisement is proprietary to Avaya to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. Enable fast advertisement if you require fast VRRP advertisement.

2. Configure VLAN 3 on router R1.

    a. Configure VLAN 3 on router R1 using spanning tree group 1.

    ```
    ERS4000# config terminal
    ERS4000# vlan create 3 type port
    ```

b. Configure the ports for VLAN 3 on R1.

```
ERS4000# config terminal
ERS4000(config)# vlan members add 3 1/14
```

c. Configure an IP address for VLAN 3.

Add IP address 10.1.21.2 / 255.255.255.0 to VLAN 3.

```
ERS4000# config terminal
ERS4000(config)# interface vlan 3
ERS4000(config)# ip address 10.1.21.2 255.255.255.0
```

d. Configure an OSPF interface for VLAN 3.

```
ERS4000# config terminal
ERS4000(config)# router ospf enable
ERS4000(config)# router ospf
ERS4000(config-router)# network 10.1.21.2
```

e. Configure VRRP on VLAN 3.

The VRRP VIP address of 10.1.21.1 is added to VLAN 2 using a VRID of 2.

```
ERS4000# config terminal
ERS4000(config)# router vrrp ena
ERS4000(config)# interface vlan 3
ERS4000(config-if)# ip vrrp address 2 10.1.21.1
ERS4000(config-if)# ip vrrp 2 priority 200
ERS4000(config-if)# ip vrrp 2 enable
```

Note:

Fast advertisement is disabled by default. Fast advertisement is proprietary to Avaya to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. Enable fast advertisement if you require fast VRRP advertisement.

3. Configure VLAN 2 on router R2.

a. Create VLAN 2 on router R2.

```
ERS4000# config terminal
ERS4000(config)# vlan create 2 type port
```

b. Configure the ports for VLAN 2 on R2.

```
ERS4000# config terminal
ERS4000(config)# vlan members add 2 1/15
```

c. Configure an IP address for VLAN 2.

Add IP address 10.1.20.3 / 255.255.255.0 to VLAN 2.

```
ERS4000# config terminal
ERS4000(config)# interface vlan 2
ERS4000(config-if)# ip address 10.1.20.3 255.255.255.0
```

d. Configure an OSPF interface for VLAN 2.

```
ERS4000# config terminal
ERS4000(config)# router ospf enable
```

```
ERS4000(config)# router ospf
ERS4000(config-router)# network 10.1.20.3
```

e. Configure VRRP on VLAN 2.

   The VRRP VIP address of 10.1.21.1 is added to VLAN 2 using a VRID of 1.

```
ERS4000# config terminal
ERS4000(config)# router vrrp ena
ERS4000(config)# interface vlan 2
ERS4000(config-if)# ip vrrp address 1 10.1.20.1
ERS4000(config-if)# ip vrrp 1 enable
ERS4000(config-if)# ip vrrp 1 priority 200
```

Note:

For this example the VRRP priority value is set to 200. This allows router R2 to be elected as the VRRP master router.

Note:

Fast advertisement is disabled by default. Fast advertisement is proprietary to Avaya to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. Enable fast advertisement if you require fast VRRP advertisement.

4. Configure VLAN 3 on router R2.

   a. Configure VLAN 3 on router R2.

```
ERS4000# config terminal
ERS4000(config)# vlan create 3 type port
```

   b. Configure the ports for VLAN 3 on R1.

```
ERS4000# config terminal
ERS4000(config)# vlan members add 3 1/14
```

   c. Configure an IP address for VLAN 3.

   Add IP address 10.1.21.3 / 255.255.255.0 to VLAN 3.

```
ERS4000# config terminal
ERS4000(config)# interface vlan 3
ERs4000(config-if)# ip address 10.1.21.3 255.255.255.0
```

   d. Configure an OSPF interface for VLAN 3.

```
ERS4000# config terminal
ERS4000(config)# router ospf enable
ERS4000(config)# router ospf
ERS4000(config-router)# network 10.1.21.3
```

   e. Configure VRRP on VLAN 3.

   The VRRP VIP address of 10.1.21.1 is added to VLAN 2 using a VRID of 2.

```
ERS4000# config terminal
ERS4000(config)# router vrrp ena
ERS4000(config)# interface vlan 3
```

```
ERS4000(config-if)# ip vrrp address 2 10.1.21.1
ERS4000(config-if)# ip vrrp 2 enable
```

Note:

Fast advertisement is disabled by default. Fast advertisement is proprietary to Avaya to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. Enable fast advertisement if you require fast VRRP advertisement.

Once you complete the VRRP configuration, use the **`show ip vrrp`** and **`show ip vrrp interface verbose`** commands to display VRRP configuration information and statistics.

VRRP configuration example 2

The figure below, ***Example VRRP topology 2***, shows two virtual routers configured on the interfaces that connect two switches to the four end hosts in the LAN.

The first virtual router is configured with a VRID of 1 and a virtual IP address of IP1. The second virtual router is configured with a VRID of 2 and a virtual IP address of IP2.

The two switches (S1 and S2) are configured with IP addresses (IP1 for S1 and IP2 for S2).

When VRRP is enabled on both switches, S2 performs as a master for VRID2 and also provides backup service for VRID1. S1 is the backup router for VRID2.

Hosts H1 and H2 are both configured with the default gateway address IP1 and hosts H3 and H4 are both configured with the default gateway address IP2.

When both switches are functioning normally, this configuration provides load splitting between S1 and S2, and full redundancy between VRID1 and VRID2.

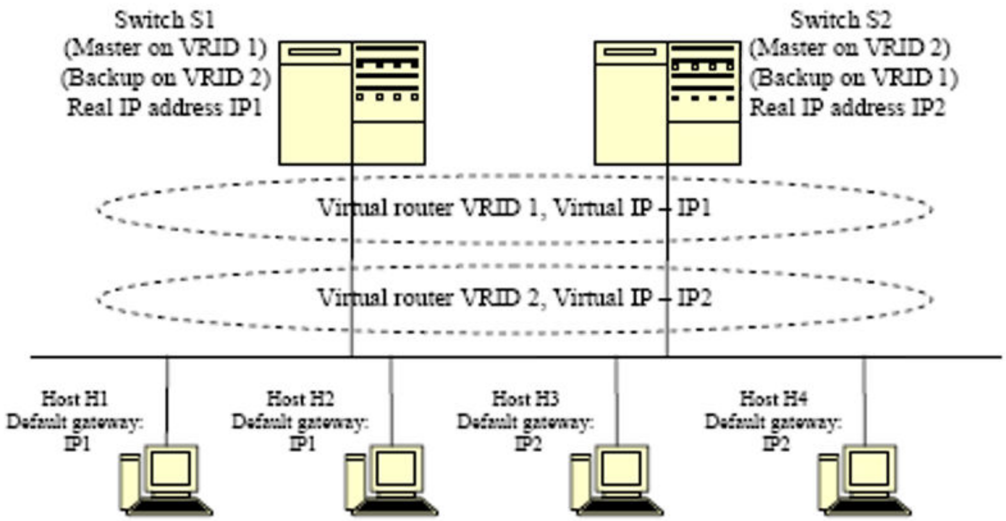

Figure 38: Example VRRP topology 2

For this configuration example, based on the VRRP topology shown above, the following apply:

- The LAN subnet is 10.1.1.0/24.
- Port 1/1 on both S1 and S2 are members of VLAN 10
- The IP address for VLAN 10 on S1 (IP1) is 10.1.1.253. This is also the default gateway address for H1 and H2.
- The IP address for VLAN 10 on S2 (IP2) is 10.1.1.254. This is also the default gateway address for H3 and H4.
- The following IP addresses are configured on the hosts:
    - H1: 10.1.1.1
    - H2:10.1.1.2
    - H3:10.1.1.3
    - H4:10.1.1.4
- VRRP is configured for S1 to back up the real IP interface on S2 and for S2 to back up the real IP interface on S1.
- VRRP licenses are available on both S1 and S2.

Configuration steps

1. Create VLAN 10 on S1 and assign an IP address.

```
S1#configure terminal
S1(config)#vlan create 10 type port
S1(config)#vlan member remove 1 1/1
S1(config)#vlan member add 10 1/1
```

```
S1(config)#interface vlan 10
S1(config-if)#ip address 10.1.1.253 255.255.255.0
```

2. Create VLAN 10 on S2 and assign an IP address.

```
S2#configure terminal
S2(config)#vlan create 10 type port
S2(config)#vlan member remove 1 1/1
S2(config)#vlan member add 10 1/1
S2(config)#interface vlan 10
S2(config-if)#ip address 10.1.1.254 255.255.255.0
```

3. On S1, configure VRID1 to back up the S1 real IP interface.

```
S1#configure terminal
S1(config)#interface vlan 10
S1(config-if)#ip vrrp address 1 10.1.1.253
S1(config-if)#ip vrrp 1 enable
```

4. Configure a virtual interface on S2, also using VRID1, to back up the real interface on S1.

```
S2#configure terminal
S2(config)#interface vlan 10
S2(config-if)#ip vrrp address 1 10.1.1.253
S2(config-if)#ip vrrp 1 enable
```

5. On S2, configure VRID2 to back up the S2 real IP interface.

```
S1#configure terminal
S1(config)#interface vlan 10
S1(config-if)#ip vrrp address 1 10.1.1.254
S1(config-if)#ip vrrp 2 enable
```

6. Configure a virtual interface on S1, also using VRID2, to back up the real interface on S2.

```
S2#configure terminal
S2(config)#interface vlan 10
S2(config-if)#ip vrrp address 1 10.1.1.254
S2(config-if)#ip vrrp 2 enable
```

7. Enable VRRP globally on S1.

```
S1#configure terminal
S1(config)#router vrrp enable
```

8. Enable VRRP globally on S2.

```
S2#configure terminal
S2(config)#router vrrp enable
```

With this configuration, S1 is the master router on IP address 10.1.1.253 and S2 is the master router on IP address 10.1.1.254. The maximum priority value of 255 is automatically configured for the interfaces on both master routers. S1 is the backup router for IP address 10.1.1.254 and S2 is the backup router for IP address 10.1.1.253. The virtual routers use the default priority of 100 for the virtual interfaces unless otherwise configured.

Chapter 12: ECMP configuration using ACLI

This section describes the procedures you can use to configure Equal Cost Multi Path (ECMP) with ACLI. With the ECMP feature routers can determine equal cost paths to the same destination prefix.

The switch can use multiple paths for traffic load sharing and in the event of network failure, faster convergence to other active paths. When the switch maximizes load sharing among equal-cost paths, the system uses links between routers more efficiently for IP traffic transmission.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Configure routing (RIP, OSPF, or static routes) on the switch.

Configuring the number of ECMP paths allotted for RIP using ACLI

Use this procedure to configure the number of ECMP paths for use by the Routing Information Protocol (RIP).

Prerequisites

- Use this command in the Global Configuration mode.

Procedure steps

Enter the following command:

```
[default] [no] rip maximum-path <path value>
```

Variable definitions

The following table describes the parameters for the rip maximum-path command.

Variable definition

Variable	Value
[default]	Resets the maximum ECMP paths allowed to the default value. DEFAULT: 1
[no]	Deletes ECMP paths specified for RIP.
<path value>	Specifies the number of ECMP paths as a value in a range from 1 to 4. DEFAULT: 1

Configuring the number of ECMP paths for OSPF using ACLI

Use this procedure to configure the number of ECMP paths for the Open Shortest Path First (OSPF) protocol.

Prerequisites

- Configure OSPF routes. See OSPF configuration using ACLI on page 97.
- Use this command in the Global Configuration mode.

Procedure steps

Enter the following command:

```
[default] [no] ospf maximum-path <path value>
```

Variable definitions

The following table describes the parameters for the **ospf maximum-path** command.

Variable definition

Variable	Value
[default]	Resets the maximum ECMP paths for OSPF to the default value. DEFAULT: 1
[no]	Deletes ECMP paths for OSPF.

Variable	Value
<path value>	Specifies the number of ECMP paths for OSPF as a value in a range from 1 to 4. DEFAULT: 1

Configuring the number of ECMP paths for static routes using ACLI

Use this procedure to configure the number of ECMP paths for static routes.

Prerequisites

- Configure static routes. See Configuring a static route on page 93.
- Use this command in the Global Configuration mode.

Procedure steps

Enter the following command:

```
[default] [no] maximum-path <path value>
```

Variable definitions

The following table describes the parameters with the **maximum-path** command.

Variable definition

Variable	Value
[default]	Resets the maximum ECMP paths for static routes to the default value. DEFAULT: 1
[no]	Deletes the ECMP paths from the static routes.
<path value>	Specifies the number of ECMP paths for static routes as a value in a range from 1 to 4. DEFAULT: 1

Displaying global ECMP path information using ACLI

Use this procedure to display ECMP path information for static routes, and RIP and OSPF protocols.

Prerequisites

- Configure the number of allowed ecmp paths for RIP/OSPF/static routes. See Configuring the number of ECMP paths for RIP/OSPF/Static on previous pages.
- Use this command in the User Exec mode.

Procedure steps

Enter the following command:

`show ecmp`

Example

Following is an example of the output for the **show ecmp** command.

```
Protocol MAX-PATH
--------    --------
static:     1
rip:        2
ospf:       4
```

ECMP configuration examples

Equal Cost Multi Path (ECMP) is an IP feature that you can use to balance routed IP traffic loads across equal-cost paths. You can use up to four equal-cost paths for each supported protocol.

ECMP supports OSPF, RIP, and static routes.

The following figure illustrates the configuration examples of ECMP:

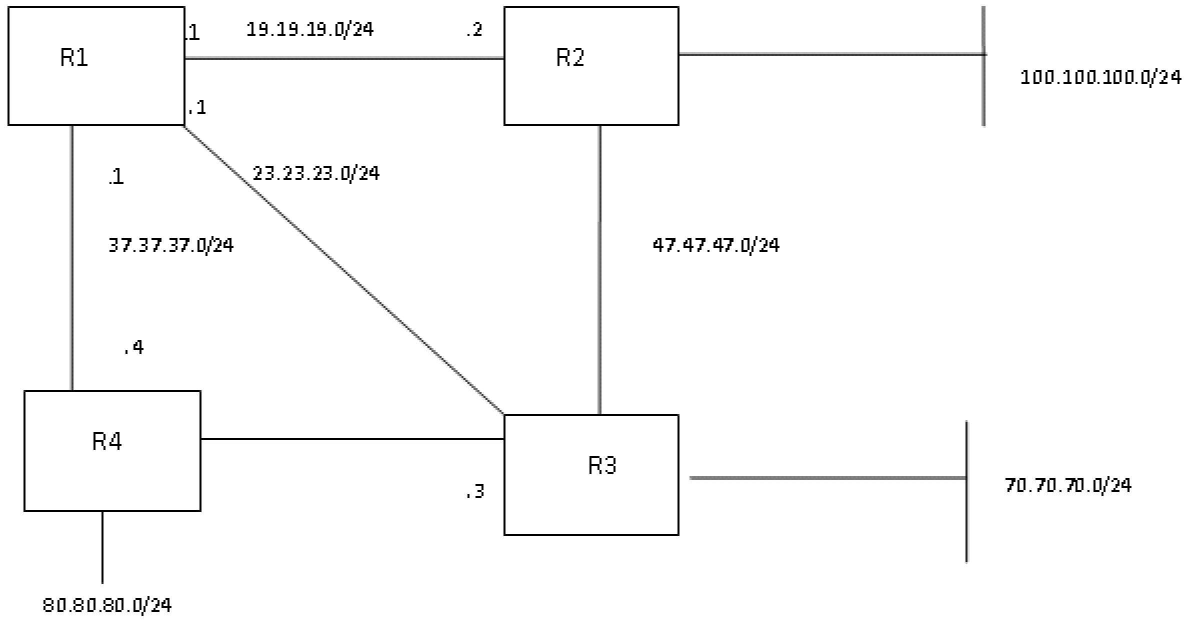

Figure 39: ECMP configuration example

Example

Consider the above setup, enable ECMP on OSPF and static routes on R1.

- `ERS4000(config)#maximum-path 2`
- `ERS4000(config)#ospf maximum-path 2`

Enable OSPF protocol on R1, R2 and R3. See OSPF configuration using ACLI on page 97 , for guidance.

Configure static routes for destination networks 80.80.80.0/24 and 100.100.100.0/24. See Configuring a static route on page 93.

- `ERS4000(config)#ip route 100.100.100.0 255.255.255.0 19.19.19.2 4`
- `ERS4000(config)#ip route 100.100.100.0 255.255.255.0 23.23.23.3 4`
- `ERS4000(config)#ip route 80.80.80.0 255.255.255.0 19.19.19.2 6`
- `ERS4000(config)#ip route 80.80.80.0 255.255.255.0 23.23.23.3 6`
- `ERS4000(config)#ip route 80.80.80.0 255.255.255.0 37.37.37.4 6`

After you complete ECMP configuration, to verify the ECMP paths in the routing table use the **`show ip route`** command.

Example

The following example displays the output for the **show ip route** command:

`ERS4000(config) #show ip route`

```
==============================================================================
Ip Route
==============================================================================
```

```
   DST          MASK            NEXT      COST  VLAN PORT PROT  TYPE   PRF
------------------------------------------------------------------------------
19.19.19.0    255.255.255.0  19.19.19.1  1     19   ----  C     DB    0
23.23.23.0    255.255.255.0  23.23.23.1  1     23   ----  C     DB    0
37.37.37.0    255.255.255.0  37.37.37.1  1     37   ----  C     DB    0
47.47.47.0    255.255.255.0  19.19.19.2  20    19   19    O     IB    20
70.70.70.0    255.255.255.0  19.19.19.2  30    19   19    O     IBE   20
                                    23.23.23.3       23   23
80.80.80.0    255.255.255.0  23.23.23.3  6     23   23    S     IBE   5
                             19.19.19.2        19   19
                             37.37.37.4        37   37
100.100.100.0 255.255.255.0  19.19.19.2  4     19   19    S     IBE   5
                             23.23.23.3        23   23
Total Routes: 11
-------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route,B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW
```

Paths shown with the letter E in the TYPE column are designated equal-cost paths.

Chapter 13: Route policies configuration using ACLI

This section describes the procedures you can use to configure route policies using ACLI.

Using standard routing schemes, packets are forwarded based on routes that have been learned by the router through routing protocols such as RIP and OSPF or through the introduction of static routes. Route policies provide the ability to forward packets based on rule sets created by the network administrator. These rule sets, or policies, are then applied to the learned or static routes.

Route policies configuration navigation

- Configuring prefix lists on page 221
- Configuring route maps on page 222
- Applying a RIP accept (in) policy on page 225
- Applying a RIP announce (out) policy on page 226
- Configuring an OSPF accept policy on page 227
- Applying the OSPF accept policy on page 228
- Displaying the OSPF accept policy on page 229
- Configuring an OSPF redistribution policy on page 229
- Applying the OSPF redistribution policy on page 231
- Displaying the OSPF redistribution policy on page 231

Configuring prefix lists

Use this procedure configure up to four prefix lists for use in route policies.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure a prefix list by using the following command:

```
[no] ip prefix-list <prefix_name> {<ip_address/mask> [ge
<mask_from>] [le <mask_to>]} [name <new_prefix_name>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
`[no]`	Removes a prefix list or a prefix from a list.
`<prefix_name>`	Specifies the name assigned to the prefix list.
`<ip_address/mask>`	Specifies the IP address and subnet mask of the prefix list. The subnet mask is expressed as a value between 0 and 32.
`ge <mask_from>`	Specifies the lower bound of the mask length. This value, when combined with the higher bound mask length (**le**), specifies a subnet range covered by the prefix list.
`le <mask_to>`	Specifies the higher bound of the mask length. This value, when combined with the lower bound mask length (**ge**), specifies a subnet range covered by the prefix list.
`name <new_prefix_name>`	Assigns a new name to previously configured prefix list.

Configuring route maps

Use this procedure to define route maps used in the configuration of route policies.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure a route map by using the following command:

```
[default] [no] route-map <map_name> [permit | deny]
<sequence_number> [enable] [match {interface <prefix_list> |
metric <metric_value> | network <prefix_list> | next-hop
<prefix_list> | protocol <protocol_name> | route-source
<prefix_list> | route-type <route_type>}] [name <new_map_name>]
[set {injectlist <prefix_list> | ip-preference <pref> | mask
<ip_address> | metric <metric_value> | metric-type
<metric_type> | nssa-pbit enable}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Configures route map default values.
[no]	Removes the specified route map.
<map_name>	Specifies the name associated with this route map.
[permit | deny]	Specifies the action to be taken when this policy is selected for a specific route. A value of permit indicates that the route is used while deny indicates that the route is ignored.
<sequence_number>	Specifies the secondary index value assigned to individual policies inside a larger policy group. Value ranges from 1 to 65535.
[enable]	Specifies whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored.
[match {interface *<prefix_list>* | metric <metric_value> |	If configured, the switch matches the specified criterion: • interface <prefix_list>—matches the IP address of the received interface against the contents of the specified prefix list. • metric <metric_value>—matches the metric of the incoming advertisement or existing route against the specified value, an

Variable	Value
network *<prefix_list>* | next-hop *<prefix_list>* | protocol *<protocol_name>* | route-source *<prefix_list>* | route-type *<route_type>*}]	integer value from 0 to 65535. If 0, then this field is ignored. The default is 0. • network <prefix_list>—matches the destination network against the contents of the specified prefix list. • next-hop <prefix_list>—matches the next hop IP address of the route against the contents of the specified prefix list. • protocol <protocol_name>—matches the protocol through which a route is learned. Options are direct, static, rip, ospf, and any. Multiple protocols can be specified by using a comma-separated list. • route-source <prefix_list>—matches the source IP address for RIP routes against the contents of the specified prefix list. • route-type <route_type>—Specifies the route type to be matched. Options are any, external, external-1, external-2, internal, and local.
[name *<new_map_name>*]	Specifies a new name to be assigned to a previously configured route map.
[set {injectlist *<prefix_list>* | ip-preference *<pref>* | mask *<mask_IP>* | metric *<metric_value>* | metric-type *<metric_type>* | nssa-pbit enable}]	If configured, the switch sets the specified parameter: • injectlist <prefix_list>: replaces the destination network of the route that matches this policy with the contents of the specified prefix list. • ip-preference <pref>: specifies the route preference value to be assigned to the route that matches this policy. Valid range is 0–255. If 0 (the default value), the global preference value is used. Used for accept policies only. • mask <mask_IP>: sets the mask IP of the route that matches this policy. Used for RIP accept policies only. • metric <metric_value>: sets the value of the metric to be assigned to matching routes. This is an integer value between 0 and 65535. • metric-type <metric_type>: sets the metric type for routes to be imported into the OSPF routing protocol. Options are type1 and type2. • nssa-pbit enable: enables the NSSA N/P-bit, which notifies the ABR to export the matching external route. Used for OSPF policies only.

Displaying route maps

Use this procedure to display configured route maps.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display route maps by using the following command:

`show route-map [detail] <map_name>`

Variable definitions

The following table describes the command variables.

Variable	Value
[detail]	Provides detailed information on the route maps.
<map_name>	Specifies the name of the route map to display.

Applying a RIP accept (in) policy

Use this procedure to specify a RIP Accept (In) policy for an interface. This policy takes the form of a previously configured route map. Only one policy can be created for each RIP interface.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Specify a RIP accept policy for an interface by using the following command:

`[default] [no] ip rip in-policy <rmap_name>`

To display RIP interface configuration, see Displaying RIP interface configuration on page 186.

Variable definitions

The following table describes the command variables.

Variable	Value
`[default]`	Removes the in policy associated with this interface.
`[no]`	Removes the in policy associated with this interface.
`<rmap_name>`	Applies the previously configured route map as the RIP accept policy.

Applying a RIP announce (out) policy

Use this procedure to specify a RIP Announce (Out) policy for an interface. This policy takes the form of a previously configured route map. Only one policy can be created for each RIP interface.

Prerequisites

- Log on to the Interface Configuration mode in ACLI.

Procedure steps

Apply a RIP Announce (Out) policy to an interface by using the following command:

```
[default] [no] ip rip out-policy <rmap_name>
```

To display RIP interface configuration, see Displaying RIP interface configuration on page 186.

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Removes the out policy associated with this interface.
[no]	Removes the out policy associated with this interface.
<rmap_name>	Applies the previously configured route map as the RIP announce policy.

Configuring an OSPF accept policy

Use this procedure to configure the router to accept advertisements from another router in the system. The referenced policy takes the form of a previously configured route map.

Accept policies are only applied to Type 5 External routes based on the advertising router ID. There can only be one OSPF accept policy on the switch and the policy is applied before updates are added to the routing table from the link state database.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure the OSPF accept-advertisements router policy by using the following command:

```
[default] [no] accept adv-rtr <router_ip_address> [enable]
[metric-type {any | type1 | type2}] [route-policy <rmap_name>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Restores an OSPF accept policy to factory defaults.
[no]	Configures the router to not accept advertisements from another router in the system.
router_ip_address	Represents the IP address of the router from which advertisements are to be accepted. The value 0.0.0.0 denotes that advertisements from all routers are accepted.
enable	Enables the accept entry for the router specified in the *<ip_address>* parameter.
metric-type {any | type1 | type2}	Indicates the type of OSPF external routes that will be accepted from this router.
route-policy *<rmap_name>*	Specifies the name of a previously configured route map to be used for filtering external routes advertised by the specified advertising router before accepting them into the routing table.

Applying the OSPF accept policy

Use this procedure to apply the configured OSPF accept policy to the switch.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Apply the OSPF accept policy to the switch by using the following command:

```
ip ospf apply accept
```

Displaying the OSPF accept policy

Use this procedure to display the OSPF accept policy.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display the OSPF accept policy by using the following command:

```
show ip ospf accept
```

Configuring an OSPF redistribution policy

Use this procedure to configure OSPF route redistribution. Redistribution of direct, RIP, and static routes is currently supported.

OSPF redistribution policies send redistributed routes as Type 5 External routes. There can be only one OSPF redistribution policy on the switch. The OSPF accept policy takes precedence over the redistribution policy.

Prerequisites

- Log on to the OSPF Router Configuration mode in ACLI.

Procedure steps

Configure OSPF route redistribution by using the following command:

```
[default] [no] redistribute <route_type> [enable] [route-policy
<rmap_name>] [metric <metric_value>] [metric-type
<metric_type>] [subnets <subnet_setting>
```

Variable definitions

The following table describes the command variables.

Variable	Value
`[default]`	Restores an OSPF route policy or OSPF route redistribution to default values.
`[no]`	Disables an OSPF route policy or OSPF route redistribution completely.
`<route_type>`	Specifies the source protocol to be redistributed. Valid options are **direct**, **rip**, and **static**.
`route-policy` *<rmap_name>*	Specifies the route policy to associate with route redistribution. This is the name of a previously configured route map.
`metric` *<metric_value>*	Specifies the metric value to associate with the route redistribution. This is an integer value between 0 and 65535.
`metric-type` *<metric_type>*	Specifies the metric type to associate with the route redistribution. Valid options are **type1** and **type2**.
`subnets` *<subnet_setting>*	Specifies the subnet advertisement setting of this route redistribution. This determines whether individual subnets are advertised. Valid options are **allow** and **suppress**.

Applying the OSPF redistribution policy

Use this procedure to apply the configured OSPF route redistribution policy to the switch.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Apply route redistribution policy to OSPF by using the following command:

```
ip ospf apply redistribute {direct | rip | static}
```

Variable definitions

The following table describes the command variables.

Variable	Value
direct	Applies only direct OSPF redistribution policy configuration to the switch
rip	Applies only RIP OSPF redistribution policy configuration on the switch.
static	Applies only static OSPF redistribution policy configuration on the switch.

Displaying the OSPF redistribution policy

Use this procedure to display the OSPF redistribution policy configuration and status.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

Display the OSPF redistribution policy by using the following command:

```
show ip ospf redistribute
```

Chapter 14: DHCP relay configuration using ACLI

This chapter describes the procedures you can use to configure Dynamic Host Configuration Protocol (DHCP) relay using the ACLI.

Prerequisites to DHCP relay configuration using ACLI

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be set as the DHCP relay agent.
- Ensure that a route (local or static) to the destination DHCP server is available on the switch.

Configuring global DHCP relay status using ACLI

Use this procedure to configure the global DHCP relay status. DHCP relay is enabled by default.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

To configure the global DHCP relay status, enter the following command:

```
[no] ip dhcp-relay
```

Variable definitions

The following table describes the **ip dhcp-relay** command variables.

Variable	Value
[no]	Disables DHCP relay.

Displaying the global DHCP relay status using ACLI

Use this procedure to display the current DHCP relay status for the switch.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

To display the global DHCP relay status, enter the following command:

```
show ip dhcp-relay
```

Specifying a local DHCP relay agent and remote DHCP server using ACLI

Use this procedure to specify a local VLAN as a DHCP relay agent on the forwarding path to a remote DHCP server. The DHCP relay agent can forward DHCP client requests from the local network to the DHCP server in the remote network.

The DHCP relay feature is enabled by default, and the default mode is BootP-DHCP.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

To configure a VLAN as a DHCP relay agent, enter the following command:

```
[no] ip dhcp-relay fwd-path <relay-agent-ip> <DHCP-server>
[enable] [disable] [mode {bootp | bootp-dhcp | dhcp}]
```

Variable definitions

The following table describes the **ip dhcp-relay fwd-path** command variables.

Variable	Value
[no]	Removes the specified DHCP forwarding path.
<relay-agent-ip>	Specifies the IP address of the VLAN that serves as the local DHCP relay agent.
<DHCP-server>	Specifies the address of the remote DHCP server to which DHCP packets are to be relayed.
[enable]	Enables the specified DHCP relay forwarding path.
[disable]	Disables the specified DHCP relay forwarding path.
[mode {bootp | bootp-dhcp | dhcp}]	Specifies the DHCP relay mode: • BootP only • BootP and DHCP • DHCP only If you do not specify a mode, the default DHCP and BootP is used.

Displaying the DHCP relay global configuration using ACLI

Use this procedure to display the current DHCP relay agent configuration for the switch.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

To display the DHCP relay configuration, enter the following command:

```
show ip dhcp-relay fwd-path
```

Job aid

The following table shows the field descriptions for the **show ip dhcp-relay fwd-path** command.

Field	Description
VLAN	Specifies the VLAN IP address.
INTERFACE	Specifies the interface IP address of the DHCP relay agent.
SERVER	Specifies the IP address of the DHCP server.
ENABLE	Specifies whether DHCP is enabled.
MODE	Specifies the DHCP mode.

Configuring the maximum packet length for DHCP relay using the ACLI

Use this procedure to configure the maximum packet length for DHCP relay.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure the maximum packet length for DHCP relay by using the following command:

`ip dhcp-relay max-frame <576-1536>`

Variable definitions

The following table defines parameters that you enter with the **`ip dhcp-relay max-frame <576-1536>`** command.

Variable	Value
`max-frame <576-1536>`	Defines the maximum DHCP relay packet length.

Configuring Option 82 for DHCP relay globally using the ACLI

Configure Option 82 for DHCP relay globally to enable or disable Option 82 for DHCP relay at the switch level.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

Configure Option 82 for DHCP relay globally by using the following command:

```
[no] ip dhcp-relay option82
```

Variable definitions

The following table defines optional parameters that you can enter with the **[no] ip dhcp-relay option82** command.

Variable	Value
[no]	Disables Option 82 for DHCP relay for the switch.

Assigning an Option 82 for DHCP Relay subscriber Id to a port using the ACLI

Assign an Option 82 for DHCP Relay subscriber Id to a port to associate an alphanumeric character string with the Option 82 function for the port.

Prerequisites

- Log on to the Ethernet Interface Configuration mode in ACLI for the port you want to modify.

Procedure steps

Assign an Option 82 for DHCP Relay subscriber Id to a port by using the following command:

```
[no] ip dhcp-relay option82-subscriber-id <WORD>
```

Variable definitions

The following table defines optional parameters that you can enter with the **[no] ip dhcp-relay option82** command.

Variable	Value
[no]	Removes the Option 82 for DHCP relay subscriber Id from a port.
<WORD>	Specifies the DHCP Option 82 subscriber Id for the port. Value is a character string between 0 and 64 characters.

Configuring DHCP relay on a VLAN using ACLI

Use this procedure to configure the DHCP relay parameters on a VLAN. To enable DHCP relay on the VLAN, enter the command with no optional parameters.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI for the VLAN on which you are configuring DHCP relay.

Procedure steps

To configure DHCP relay on a VLAN, enter the following command:

```
[no] ip dhcp-relay [broadcast] [min-sec <min-sec>] [mode {bootp
| dhcp | bootp_dhcp}] [option82]
```

Variable definitions

The following table describes the **ip dhcp-relay** command variables.

Variable	Value
[no]	Disables DHCP relay status and parameters on the specified VLAN.
[broadcast]	Enables the broadcast of DHCP reply packets to the DHCP clients on this VLAN interface.
min-sec *<min-sec>*	Indicates the min-sec value. The switch immediately forwards a BootP/DHCP packet if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped. Range is 0-65535. The default is 0.

Variable	Value
mode {bootp | dhcp | bootp_dhcp}	Specifies the type of DHCP packets this VLAN supports: • bootp - Supports BootP only • dhcp - Supports DHCP only • bootp_dhcp - Supports both BootP and DHCP
option82	Enables Option 82 for DHCP relay on a VLAN.

Displaying the DHCP relay configuration for a VLAN using ACLI

Use this procedure to display the current DHCP relay parameters configured for a VLAN.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To display the DHCP relay VLAN parameters, enter the following command:

```
show vlan dhcp-relay [<vid>]
```

Variable definitions

The following table describes the **show vlan dhcp-relay** command variables.

Variable	Value
[<vid>]	Specifies the VLAN ID of the VLAN to be displayed. Range is 1-4094.

Job aid

The following table shows the field descriptions for the **`show vlan dhcp-relay`** command.

Field	Description
IfIndex	Indicates the VLAN interface index.
MIN_SEC	Indicates the min-sec value. The switch immediately forwards a BootP/DHCP packet if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped.
ENABLED	Indicates whether DHCP relay is enabled on the VLAN.
MODE	Indicates the type of DHCP packets this interface supports. Options include none, BootP, DHCP, and both.
ALWAYS_BROADCAST	Indicates whether DHCP reply packets are broadcast to the DHCP client on this VLAN interface.
OPTION_82	Indicates if Option 82 for DHCP Relay is enabled or disabled on the VLAN.

Displaying the DHCP relay configuration for a port using ACLI

Use this procedure to display the current DHCP relay parameters configured for an Ethernet interface port.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

To display the DHCP relay port parameters, enter the following command:

```
show ip dhcp-relay interface Ethernet <slot/port>
```

Variable definitions

The following table defines parameters that you enter with the **show ip dhcp-relay interface Ethernet <slot/port>** command.

Variable	Value
<slot/port>	Specifies the slot and port number of the port to be displayed.

Displaying DHCP relay counters using ACLI

Use this procedure to display the current DHCP relay counters. This includes the number of requests and the number of replies.

Prerequisites

- Log on to the User EXEC mode in ACLI.

Procedure steps

To display the DHCP relay counters, enter the following command:

```
show ip dhcp-relay counters
```

Job aid

The following table shows the field descriptions for the **show ip dhcp-relay counters** command.

Field	Description
INTERFACE	Indicates the interface IP address of the DHCP relay agent.
REQUESTS	Indicates the number of DHCP requests.
REPLIES	Indicates the number of DHCP replies.

Clearing DHCP relay counters for a VLAN using ACLI

Use this procedure to clear the DHCP relay counters for a VLAN.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

To clear the DHCP relay counters, enter the following command:

```
ip dhcp-relay clear-counters
```

Chapter 15: UDP broadcast forwarding configuration using ACLI

This chapter describes the procedures you can use to configure UDP broadcast forwarding using ACLI. UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address.

You cannot enable or disable the UDP broadcast forwarding feature on a global level. When you attach the first UDP forwarding list to a VLAN interface, the feature is enabled. When you remove the last UDP forwarding list from a VLAN, the feature is disabled.

Important:

UDP broadcast forwarding shares resources with the Quality of Service (QoS) feature. When UDP forwarding is enabled, the switch dynamically assigns the highest available precedence value to the UDP forwarding feature. To display the assigned precedence after you enable UDP forwarding, enter the **`show qos diag`** command.

For further information on QoS policies, see *Configuring Quality of Service on Avaya Ethernet Routing Switch 4000 Series*, NN47205-504.

Prerequisites to UDP broadcast forwarding using ACLI

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a UDP forwarding interface.
- Ensure that a route (local or static) to the destination address is available on the switch.

Important:

If you configure EAPOL on the switch, enable EAPOL before enabling UDP Forwarding, otherwise the UDP broadcast traffic matching UDP forward lists is forwarded regardless of the EAPOL port state (authorized, force unauthorized, or auto).

UDP broadcast forwarding configuration procedures

To configure UDP broadcast forwarding, perform the following steps:

1. Create UDP protocol entries that specify the protocol associated with each UDP port that you want to forward.
2. Create a UDP forwarding list that specifies the destination IP addresses for each forwarding UDP port. (You can create up to 128 UDP forwarding lists.)
3. Apply UDP forwarding lists to local VLAN interfaces.

Configuring UDP protocol table entries

Use this procedure to create UDP protocol table entries that identify the protocols associated with specific UDP ports that you want to forward.

Procedure steps

To configure a UDP table entry, enter the following from the Global Configuration mode:

`ip forward-protocol udp [<forwarding_port> <protocol_name>]`

Variable definitions

The following table describes the **ip forward-protocol udp** command variables.

Variable	Value
<forwarding_port>	Specifies the UDP port number. Range is 1-65535.
<protocol_name>	Specifies the UDP protocol name.

Displaying the UDP protocol table

Use this procedure to display the configured UDP protocol table entries.

Procedure steps

To display the UDP protocol table, enter the following from the User Exec mode:

```
show ip forward-protocol udp
```

Job aid

The following table shows the field descriptions for the **show ip forward-protocol udp** command.

Field	Description
UDP_PORT	Indicates the UDP ports.
PROTOCOL_NAME	Indicates the name of the associated protocol.

Configuring a UDP forwarding list

Use this procedure to configure a UDP forwarding list, which associates UDP forwarding ports with destination IP addresses. Each forwarding list can contain multiple port/destination entries. You can configure a maximum of 16 port/destination entries in one forwarding list.

You can configure up to 128 forwarding lists.

Procedure steps

To configure a UDP port forwarding list, enter the following from the Global Configuration mode:

```
ip forward-protocol udp portfwdlist <forward_list> <udp_port>
<dest_ip> [name <list_name>]
```

Variable definitions

The following table describes the **ip forward-protocol udp portfwdlist** command variables.

Variable	Value
<forward_list>	Specifies the ID of the UDP forwarding list. Range is 1-128.
<udp_port>	Specifies the port on which the UDP forwarding originates.
<dest_ip>	Specifies the destination IP address for the UDP port.

Variable	Value
<list_name>	Specifies the name of the UDP forwarding list being created (maximum 15 characters).

Applying a UDP forwarding list to a VLAN

Use this procedure to associate a UDP forwarding list with a VLAN interface (you can attach only one list at a time to a VLAN interface).

You can bind the same UDP forwarding list to a maximum of 16 different VLANs.

Procedure steps

To associate a UDP forwarding list to a VLAN, enter the following from the VLAN Interface Configuration mode:

```
ip forward-protocol udp [vlan <vid>] [portfwdlist
<forward_list>] [broadcastmask <bcast_mask>] [maxttl <max_ttl>]
```

Variable definitions

The following table describes the **ip forward-protocol udp** command variables.

Variable	Value
<vid>	Specifies the VLAN ID on which to attach the UDP forwarding list. This parameter is optional, and if not specified, the UDP forwarding list is applied to the interface specified in the **interface vlan** command.
<forward_list>	Specifies the ID of the UDP forwarding list to attach to the selected VLAN interface.
<bcast_mask>	Specifies the 32-bit mask used by the selected VLAN interface to make forwarding decisions based on the destination IP address of the incoming UDP broadcast traffic. If you do not specify a broadcast mask value, the switch uses the mask of the interface to which the forwarding list is attached. (See Note 1.)
<max_ttl>	Specifies the timet-to-live (TTL) value inserted in the IP headers of the forwarded UDP packets coming out of the selected VLAN interface. If you do not specify a TTL value, the default value (4) is used. (See Note 1.)

Variable	Value
Note 1: If you specify maxttl and/or broadcastmask values with no portfwdlist specified, the switch saves the settings for this interface. If you subsequently attach portfwdlist to this interface without defining the maxttl and/or broadcastmask values, the saved parameters are automatically attached to the list. But, if when specifying the portfwdlist, you also specify the maxttl and/or broadcastmask, your specified properties are used, regardless of any previous configurations.	

Displaying the UDP broadcast forwarding configuration

Use this procedure to display the UDP broadcast forwarding configuration.

Procedure steps

To display the UDP broadcast forwarding configuration, enter the following from the User Exec mode:

```
show ip forward-protocol udp [interface [vlan <1-4094>]]
[portfwdlist [<portlist>]
```

Variable definitions

The following table describes the **show ip forward-protocol udp** command variables.

Variable	Value
[interface [vlan *<1-4094>*]]	Displays the configuration and statistics for a VLAN interface. If no VLAN is specified, the configuration for all UDP forwarding-enabled VLANs is displayed.
[portfwdlist [*<forward_list>*]	Displays the specified UDP forwarding list. If no list is specified, a summary of all forwarding lists is displayed.

Job aids

The following table shows the field descriptions for the **show ip forward-protocol udp** command.

Field	Description
UDP_PORT	Indicates the UDP ports.
PROTOCOL_NAME	Indicates the name of the protocol.

The following table shows the field descriptions for the **show ip forward-protocol udp interfaces** command.

Field	Description
INTF_ADDR	Indicates the IP address of the interface.
FWD LISTID	Identifies the UDP forwarding policy.
MAXTTL	Indicates the maximum TTL.
RXPKTS	Indicates the number of received packets.
FWDPKTS	Indicates the number of forwarded packets.
DRPDEST UNREACH	Indicates the number of dropped packets that cannot reach the destination.
DRP_UNKNOWN PROTOCOL	Indicates the number of packets dropped with an unknown protocol.
BDCASTMASK	Indicates the value of the broadcast mask.

The following table shows the field descriptions for the **show ip forward-protocol udp portfwdlist** command.

Field	Description
LIST_ID	Specifies the UDP forwarding policy number.
NAME	Specifies the name of the UDP forwarding policy.

Clearing UDP broadcast counters on an interface

Use this procedure to clear the UDP broadcast counters on an interface.

Procedure steps

To clear the UDP broadcast counters, enter the following from the Privileged Exec command mode:

```
clear ip forward-protocol udp counters <1-4094>
```

Variable definitions

The following table describes the **clear ip forward-protocol udp counters** command variables.

Variable	Value
<1-4094>	Specifies the VLAN ID.

Chapter 16: Directed broadcasts configuration using ACLI

This chapter describes the procedures you can use to configure and display the status of directed broadcasts using ACLI.

Configuring directed broadcasts

Use this procedure to enable directed broadcasts on the switch. By default, directed broadcasts are disabled.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a broadcast interface.
- Ensure that a route (local or static) to the destination address is available on the switch.

Procedure steps

To enable directed broadcasts, enter the following from the Global Configuration mode:

```
ip directed-broadcast enable
```

Displaying the directed broadcast configuration

Use this procedure to display the status of directed broadcasts on the switch. By default, directed broadcasts are disabled.

Procedure steps

To display directed broadcast status, enter the following from the User EXEC mode:

```
show ip directed-broadcast
```

Chapter 17: Static ARP and Proxy ARP configuration using ACLI

This chapter describes the procedures you can use to configure Static ARP, Proxy ARP, and display ARP entries using the ACLI.

Static ARP configuration

This section describes how to configure Static ARP using the ACLI.

Configuring a static ARP entry

Use this procedure to configure a static ARP entry.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN.
- Use this command in the Global Configuration mode.

Procedure steps

Enter the following command:

```
[no] arp <A.B.C.D> <aa:bb:cc:dd:ee:ff> <unit / port> [id <1-4094>]
```

Configuration example: Adding a static ARP entry to a VLAN

The following is an example of adding a static ARP entry to a VLAN or brouter port:

Example

```
ERS4000(config)#arp 10.1.1.23 00:00:11:43:54:23 1/48 id 1
```

Variable definitions

The following table describes the parameters for the **arp** command.

Variable	Value
[no]	Removes the specified ARP entry.
<A.B.C.D>	Specifies the IP address of the device being set as a static ARP entry.
<aa:bb:cc:dd:ee:ff>	Specifies the MAC address of the device being set as a static ARP entry.
<unit / port>	Specifies the unit and port number to which the static ARP entry is being added.
id *<1 - 4094>*	Specifies the VLAN ID to which the static ARP entry is being added.

Displaying the ARP table

Use the following procedures to display the ARP table, configure a global timeout for ARP entries, and clear the ARP cache.

Displaying ARP entries

Use this procedure to display ARP entries.

Prerequisites

Use this command in the User EXEC mode. The **show arp** command is invalid if the switch is not in Layer 3 mode

Procedure steps

Enter the following command:

```
show arp-table
```

OR

```
show arp [<ip-addr>] [-s <subnet> <mask>] [static <ip-addr> [-
s <subnet> <mask>]] [dynamic <ip-addr> [-s <subnet> <mask>]]
[summary]
```

Variable definitions

The following table describes the parameters for the **show arp** command.

Variable	Value
`<ip-addr>`	Specifies the IP address of the ARP entry to be displayed.
`-s` *<subnet> <mask>*	Displays ARP entries for the specified subnet only.
`static` *<ip-addr>* `[-s` *<subnet> <mask>*`]`	Displays static entries for the specified subnet. If you do not specify a subnet, all configured static entries are displayed, including those without a valid route.
`dynamic` *<ip-addr>* `[-s` *<subnet> <mask>*`]`	Displays dynamic entries for the specified subnet. If you do not specify a subnet, all dynamic entries are displayed.
`summary`	Displays a summary of ARP entries.

Job aid

The following table shows the field descriptions for **show arp-table** and **show ip arp** commands.

Field	Description
IP Address	Specifies the IP address of the ARP entry.
Age (min)	Displays the ARP age time.
MAC Address	Specifies the MAC address of the ARP entry.
VLAN-Unit/Port/Trunk	Specifies the VLAN/port of the ARP entry.
Flags	Specifies the type of ARP entry: S=Static, D=Dynamic, L=Local, B=Broadcast.

Configuring a global timeout for ARP entries

Use this procedure to configure an aging time for the ARP entries.

Prerequisites

Use this command in the Global Configuration mode.

Procedure steps

Enter the following command:

```
arp timeout <5-360>
```

Variable definitions

The following table describes the parameters for **arp timeout** command.

Variable	Value
<5-360>	Specifies the amount of time in minutes before an ARP entry ages out. The range is 5 to 360 minutes. DEFAULT: 360 minutes

Restoring default timeout for ARP entries using ACLI

Use this procedure to return the aging time for the ARP entries to the default value.

Prerequisites

- Use this command in the Global Configuration mode.

Procedure steps

Use this command in the Global Configuration mode.

```
default arp timeout
```

Variable definitions

The following table describes the parameters for the **arp timeout** command.

Variable definition

Variable	Value
default	Returns the amount in time in seconds before an ARP entry ages out to the default value. DEFAULT: 21600 seconds

Clearing the ARP cache

Use this procedure to clear the cache of ARP entries.

Procedure steps

To clear the ARP cache, enter the following from the Global Configuration mode:

```
clear arp-cache
```

Proxy ARP configuration

This section describes how to configure Proxy ARP using the ACLI.

Configuring proxy ARP status

Use this procedure to enable proxy ARP functionality on a VLAN. By default, proxy ARP is disabled.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a Proxy ARP interface.

Procedure steps

To configure proxy ARP status on a VLAN, enter the following from the VLAN Interface Configuration mode:

```
[default] [no] ip arp-proxy enable
```

Variable definitions

The following table describes the **ip arp-proxy enable** command variables.

Variable	Value
[default]	Disables proxy ARP functionality on the VLAN.
[no]	Disables proxy ARP functionality on the VLAN.

Displaying proxy ARP status on a VLAN

Use this procedure to display the status of proxy ARP on a VLAN.

Procedure steps

To display proxy ARP status for a VLAN, enter the following from the User EXEC mode:

`show ip arp-proxy interface [vlan <vid>]`

Variable definitions

The following table describes the **show ip arp-proxy interface** command variables.

Variable	Value
<vid>	Specifies the ID of the VLAN to display. Range is 1-4094.

Job aid

The following table shows the field descriptions for the **show ip arp-proxy interface** command.

Field	Description
Vlan	Identifies a VLAN.
Proxy ARP status	Specifies the status of Proxy ARP on the VLAN.

Chapter 18: IP blocking configuration using ACLI

This chapter describes the procedures you can use to configure and display the status of IP blocking in a stack using ACLI.

Configuring IP blocking for a stack

Use this procedure to set the IP blocking mode in the stack.

Procedure steps

To configure IP blocking, enter the following from the Global Configuration mode:

```
ip blocking-mode {full | none}
```

Variable definitions

The following table describes the **ip blocking-mode** command variables.

Variable	Value
full	Select this parameter to set IP blocking to full, which never allows a duplicate IP address in a stack.
none	Select this parameter to set IP blocking to none, which allows duplicate IP addresses unconditionally.

Displaying IP blocking status

Use this command to display the status of IP blocking on the switch.

Procedure steps

1. To display the IP blocking mode on the switch, enter the following from the User EXEC mode:

   ```
   show ip blocking-mode
   ```

2. To display the IP blocking state on the switch, enter the following from the User EXEC mode:

   ```
   show ip-blocking
   ```

Chapter 19: IGMP snooping configuration using ACLI

This chapter describes the procedures you can use to configure and display IGMP snooping parameters using ACLI.

Note:

As of Release 5.6, many of the vlan igmp commands have been superseded by the newer ip igmp commands. The vlan igmp commands in these cases are maintained for backwards compatibility only.

Displaying the switch IGMP snooping configuration status using ACLI

Use this procedure to display information about the IGMP snooping configuration for the switch.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp snooping
```

Job Aid

The following table describes the information displayed with the **`show ip igmp snooping`** command.

`Vlan`	Indicates the VLAN ID
`Snoop Enable`	Indicates whether snoop is enables (true) or disabled (false)
`Proxy Snoop Enable`	Indicates whether IGMP proxy is enabled (true) or disabled (false)

`Static Mrouter Ports`	Indicates the static mrouter ports in this VLAN that provide connectivity to an IP multicast router.
`Active Mrouter Ports`	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
`Mrouter Expiration Time`	Specifies the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Displaying IGMP interface information using ACLI

Use this procedure to display configuration information for all IGMP interfaces, or for a specific VLAN.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp interface [vlan <vid>]
```

Variable definitions

Variable	Value
vlan <vid>	Specifies a specific VLAN for which to display IGMP interface information.

Job aid

The following table shows the field descriptions for **show ip igmp interface** command output.

Field	Description
VLAN	Indicates the VLAN on which IGMP is configured.
Query Intvl	Specifies the frequency (in seconds) at which host query packets are transmitted on the interface.
Vers	Specifies the version of IGMP configured on this interface.
Oper Vers	Specifies the version of IGMP running on this interface.
Querier	Specifies the IP 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 Interface version. You must configure all routers on a LAN to run the same version of IGMP. Thus, if queries are received with the wrong version, a configuration error occurs.
Joins	Indicates the number of times a group membership was added on this interface.
Robust	Specifies the robust value configured for expected packet loss on the interface.
LastMbr 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 the interface is configured for IGMPv1.

Field	Description
Send Query	Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES or NO. Default is disabled.

Creating an IGMP VLAN interface using ACLI

Use this procedure to create a new IGMP interface.

Important:

You can create a maximum of 256 IGMP VLAN interfaces.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
ip igmp
```

Deleting an IGMP VLAN interface using ACLI

Use this procedure to remove an IGMP interface.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter either of the following commands:

```
default ip igmp
```

OR

```
no ip igmp
```

Note:

When you remove an IGMP VLAN interface, the system restores the default values of any previously saved IGMP parameters (e.g.. snooping, proxy, mrouter, robust-value, etc).

Enabling or disabling IGMP snooping for a VLAN using ACLI

Use the following procedure to enable or disable IGMP snooping for a VLAN.

IGMP snooping is disabled by default.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] [no] ip igmp snooping
```

Variable definitions

Variable	Value
`default`	Disables IGMP snooping on the selected VLAN.
`no`	Disables IGMP snooping on the selected VLAN.

Adding static mrouter ports to a VLAN using ACLI

IGMP snoop considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port. By default, the switch forwards incoming IGMP membership reports only to the active mrouter port.

To forward the IGMP reports to additional ports, you can configure the additional ports as static mrouter ports.

Important:

The static mrouter port version must match the IGMP version configured on the VLAN of the IGMP querier router.

Use this procedure to add one or more static mrouter ports to a VLAN (IGMPv1, IGMPv2, and IGMPv3 according to the supported version on the VLAN).

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
ip igmp mrouter <port_list>
```

Variable definitions

Variable	Value
<port_list>	Specifies the port or ports to add to the VLAN as static mrouter ports.

Removing static mrouter ports from a VLAN using ACLI

Use this procedure to remove one or more static mrouter ports from a VLAN.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

1. At the VLAN Interface Configuration prompt, enter either of the following commands to remove all static mrouter ports from the VLAN:

   ```
   default ip igmp mrouter
   ```

   OR

   ```
   no ip igmp mrouter
   ```

2. At the VLAN Interface Configuration prompt, enter either of the following commands to remove specific static mrouter ports from the VLAN:

   ```
   default ip igmp mrouter
   ```

   OR

   ```
   no ip igmp mrouter <port_list>
   ```

Variable definitions

Variable	Value
<port_list>	Specifies the static mrouter port or ports to remove from the VLAN.

Enabling or disabling IGMP proxy on a VLAN using ACLI

Use this procedure to enable or disable IGMP proxy on a VLAN. When IGMP proxy is enabled, the switch consolidates incoming report messages into one proxy report for that group.

If IGMP snooping was not previously enabled on A VLAN, snooping is enabled automatically when you enable IGMP proxy on that VLAN.

IGMP proxy is disabled by default.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] [no] ip igmp proxy
```

Variable definitions

Variable	Value
default	Disables IGMP proxy on the selected VLAN.
no	Disables IGMP proxy on the selected VLAN.

Configuring IGMP snooping robustness for a VLAN using ACLI

Use this procedure to set the robustness value for a VLAN. With IGMP snooping robustness, the switch can offset expected packet loss on a subnet.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] ip igmp robust-value <2-255>
```

Variable definitions

Variable	Value
default	Sets the IGMP snooping robustness to the default value of 2.
<2–255>	Specifies a numerical value for IGMP snooping robustness. Values range from 2 to 255.

Configuring the IGMP last member query interval for a VLAN using ACLI

Use this procedure to set the maximum response time (in tenths of a second) that is inserted into group-specific queries that are sent in response to leave group messages.

IGMP also uses the last member query interval as the period between group specific query messages.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the the VLAN Interface Configuration prompt, enter the following command:

```
[default] ip igmp last-member-query-interval <0-255>
```

Variable definitions

Variable	Value
<0–255>	Specifies the last member query interval value in 1/10 of a second. Values range from 0 to 255. Avaya recommends that you configure this parameter to values higher than 3. If a fast leave process is not required, Avaya recommends values above 10.
[default]	Sets the last member query interval to the default value of 10.

Configuring the IGMP query interval for a VLAN using ACLI

Use this procedure to set the frequency (in seconds) at which host query packets are transmitted on the VLAN.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] ip igmp query-interval <1-65535>
```

Variable definitions

Variable	Value
<1–65535>	Specifies the query interval value. Values range from 1 to 65535 seconds.
[default]	Sets the query interval to the default value of 125 seconds.

Configuring the IGMP maximum query response time for a VLAN using ACLI

Use this procedure to set the maximum response time (in tenths of a second) that is advertised in IGMPv2 general queries on the VLAN.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] ip igmp query-max-response <0-255>
```

Variable definitions

Variable	Value
[default]	Sets the maximum query response time to the default value of 100.
<0–255>	Specifies the maximum query response time value in 1/10 of a second. Values range from 0 to 255.

Enabling or disabling IGMP send query on a VLAN using ACLI

Use this procedure to enable or disable IGMP send query on a snoop-enabled VLAN.

When IGMP send query is enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switch or host that wants to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate multicast group packet forwarding.

IGMP send query is disabled by default.

Prerequisites

- Enable IGMP snooping on the VLAN
- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] [no] ip igmp send-query
```

Variable definitions

Variable	Value
default	Disables IGMP send query on the selected VLAN.
no	Disables IGMP send query on the selected VLAN.

Configuring the IGMP version on a VLAN using ACLI

Use this procedure to configure the IGMP version running on the VLAN. You can specifiy the version as IGMPv1, IGMPv2, or IGMPv3.

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the the VLAN Interface Configuration, enter the following command:

```
[default] ip igmp version <1-3>
```

Variable definitions

Variable	Value
default	Restores the IGMP protocol version to the default value (IGMPv2).
<1-3>	Specifies the IGMP version. • 1—IGMPv1 • 2—IGMPv2 • 3—IGMPv3

Enabling or disabling IGMP router alert on a VLAN using ACLI

Use this procedure to enable the router alert feature. This feature instructs the router to drop control packets that do not have the router-alert flag in the IP header.

Important:

To maximize your network performance, Avaya recommends that you set the router alert option according to the version of IGMP currently in use:

- IGMPv1—Disable
- IGMPv2—Enable
- IGMPv3—Enable

Prerequisites

- Log on to the VLAN Interface Configuration mode in ACLI.

Procedure steps

At the VLAN Interface Configuration prompt, enter the following command:

```
[default] [no] ip igmp router-alert
```

Variable definitions

Variables	Description
`default`	Disables the router alert option.
`no`	Disables the router alert option.

Displaying IGMP router alert configuration information using ACLI

Use this procedure to display configuration information for the IGMP router alert feature.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp router-alert [vlan <1-4094>]
```

Variable definitions

Variable	Value
vlan *<1-4094>*	Displays IGMP router alert configuration information for a specific VLAN. • *<1-4094>*—specifies the VLAN ID.

Applying the IGMP filter profile on an Ethernet interface using ACLI

Use this procedure to apply the IGMP filter profile on an Ethernet interface.

Prerequisites

- Log on to the Ethernet Interface Configuration mode in ACLI.

Procedure steps

At the Ethernet interface configuration prompt, enter the following command:

```
ip igmp filter <1-65535>
```

Variable definitions

Variable	Value
<1-65535>	Specifes a profile ID. Values range from 1 to 65535.

Clearing IGMP profile statistics using ACLI

Use this procedure to clear IGMP statistics for a selected profile, or all profiles.

Prerequisites

- Use this command in the Privileged Exec mode.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
clear ip igmp profile stats [<1-65535>]
```

Variable definitions

Variable	Value
<1-65535>	Specifies the profile ID. If you do not include this variable in the command, statistics for all profiles are cleared.

Enabling or disabling unknown multicast flooding using ACLI

Use this procedure to enable or disable the functionality for the switch to flood all VLANs with unknown multicast addresses.

Unknown multicast flooding is disabled by default.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter the following command:

```
[default] vlan igmp unknown-mcast-no-flood {enable | disable}
```

Variable definitions

Variable	Value
`default`	Enables the flooding of multicast packets on all VLANs.
`enable`	Prevents the flooding of multicast packets on all VLANs.
`disable`	Enables the flooding of multicast packets on all VLANs.

Displaying IGMP profiles using ACLI

Use this procedure to display information for a specific IGMP profile or for all IGMP profiles configured on the switch.

Prerequisites

- Log on to the Privileged Exec mode in ACLI.

Procedure steps

At the Privileged Exec prompt, enter the following command:

```
show ip igmp profile [<1-65535>]
```

Variable definitions

Variables	Description
<1-65535>	Specifies a profile ID. Values range from 1 to 65535.

Configuring an IGMP profile using ACLI

Use this procedure to create an IGMP profile and set the profile range start and end IP addresses for the new profile, or to set the profile range start and end IP addresses for an existing IGMP profile.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

1. At the Global Configuration prompt, enter the following command to create a new profile or access an existing profile:

   ```
   ip igmp profile <1-65535>
   ```

2. At the `config-igmp-profile` prompt, enter the following command:

   ```
   range <start_ip_address> <end_ip_address>
   ```

Variable definitions

Variables	Description
<1-65535>	Specifies a profile ID. Values range from 1 to 65535.
<start_ip_address>	Specifies the first IP address in the IGMP profile range, in the **A.B.C.D** format.
<end_ip_address>	Specifies the last IP address in the IGMP profile range, in the **A.B.C.D** format.

Deleting an IGMP profile using ACLI

Use this procedure to remove an IGMP profile and the IP address range configured for that profile, from the switch.

Prerequisites

• Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter either of the following commands:

```
no ip igmp profile <1-65535>
```

OR

```
default ip igmp profile <1-65535>
```

Variable definitions

Variables	Description
<1-65535>	Specifies a profile ID. Values range from 1 to 65535.

Enabling or disabling an IGMP profile on a port

Use this procedure to add or remove an IGMP profile on an interface port.

Procedure

1. Enter Ethernet Interface Configuration mode:
   ```
   enable
   configure terminal
   interface ethernet <port>
   ```
2. To add an IGMP profile on the port, enter the following command:
   ```
   ip igmp filter <1-65535>
   ```
3. To remove an IGMP profile on the port, enter the following command:

```
no ip igmp filter <1-65535>
```

Variable definitions

The following table describes the parameters for the **ip igmp filter** command.

Variables	Description
<1–65535>	Specifies an IGMP profile ID. Values range from 1 to 65535.

Displaying IGMP cache information using ACLI

Use this procedure to display the learned multicast groups in the cache and the IGMPv1 version timers.

Note:

Using the **show ip igmp cache** command may not display the expected results in some configurations. If the expected results are not displayed, use the **show ip igmp group** command to view the information.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp cache
```

Job Aid

The following table shows the field descriptions for the **show ip igmp cache** command.

Field	Description
Group Address	Indicates the multicast group address
VLAN ID	Indicates the VLAN interface on which the group exists.
Last Reporter	Indicates the last IGMP host to join the group.
Expiration	Indicates the group expiration time (in seconds).
V1 Host Timer	Indicates the time remaining until the local router assumes that no IGMP version 1 members exist on the IP subnet attached to the interface. Upon hearing an IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local interface ignores IGMPv2 Leave messages that it receives for this group.
Type	Indicates whether the entry is learned dynamically or is added statically.

Displaying IGMP group information using ACLI

Use this procedure to display the IGMP group information for the switch.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp group [count] [group <A.B.C.D>] [member-subnet
<A.B.C.D>/<0-32>
```

Variable definitions

Variable	Value
`count`	Displays the number of IGMP group entries.
group <A.B.C.D>	Displays group information for the specified group.
member-subnet<A.B.C.D> / <0–32>	Displays group information for the specified member subnet.

Job aid

The following table shows the field descriptions for the **`show ip igmp group`** command.

Field	Description
Group Address	Indicates the multicast group address
VLAN	Indicates the VLAN interface on which the group exists.
Member Address	Indicates the IP address of the IGMP receiver (host or IGMP reporter). The IP address is 0.0.0.0 if the type is static.
Expiration	Indicates the time left before the group report expires. This variable is updated upon receiving a group report.
Type	Specifies the type of membership: static or dynamic
In Port	Identifies the member port for the group. This is the port on which group traffic is forwarded and in those case where the type is dynamic, it is the port on which the IGMP join was received.

Displaying extended IGMP group information using ACLI

Use this procedure to return all the information returned by the `show ip igmp group` command and some extended information useful only in an IGMP V3 environment.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

Enter the following command:

```
show ip igmp group-ext [count] [group <A.B.C.D>] [member-subnet
<A.B.C.D/0-32>] [source < A.B.C.D>]
```

Variable definitions

Variable	Value
count	Displays the entry count for IGMP group extended details.
group<*A.B.C.D*>	Displays IGMP group extended details for the selected group. • A.B.C.D—specifies the group IP address.
member-subnet<*A.B.C.D/0-32*>	Displays IGMP group extended details for the selected member subnet. • A.B.C.D—specifies the member IP address. • 0-32—specifies the subnet for the member IP address.
source< *A.B.C.D*>	Displays IGMP group extended details for the selected source IP address. • A.B.C.D—specifies the source IP address.

Job aid

The output of the **`show ip igmp group-ext`** command includes all the information returned by the **`show ip igmp group`** command and has two additional fields:

- Source Address: indicates the source address specified in the Source Address field of the group record(s) of IGMP V3 reports. A separate entry is returned for every source address registered.
- Mode: shows the group record type of the IGMP entry.

The following example displays sample output for the **show ip igmp group-ext count** command:

```
ERS-4524GT#show ip igmp group-ext count
Igmp Group (Receiver) Count:  0
Igmp Entry Available:  512
```

Displaying VLAN multicast address flooding information using ACLI

Use this procedure to display information about MAC and IP addresses that are configured to flood VLANs with unknown multicast packets.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show vlan igmp unknown-mcast-allow-flood [<vid>]
```

Variable definitions

Variable	Value
<vid>	Specifies a specific VLAN identifier for which to display multicast flooding information.

Configuring multicast VLAN flooding using ACLI

Use this procedure to specify MAC addresses and IP addresses to flood VLANs with unknown multicast packets.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter the following command:

```
[default] [no] vlan igmp unknown-mcast-allow-flood <vid>
<A.B.C.D> <H.H.H>
```

Variable definitions

Variable	Value
default	Deletes the selected MAC or IP addresses from the respective MAC or IP multicast filter table.
no	Deletes the selected MAC or IP addresses from the respective MAC or IP multicast filter table.
<A.B.C.D>	Specifies the multicast IP address to flood VLANs.
<H.H.H>	Specifies the multicast MAC address to flood VLANs. Accepted formats are: • H.H.H • xx:xx:xx:xx:xx:xx • xx.xx.xx.xx.xx.xx • xx-xx-xx-xx-xx-xx
<vid>	Specifies the VLAN ID.

Job aid

The following table shows the field descriptions for the **`show vlan multicast membership`** command.

Field	Description
Multicast Group Address	Indicates the multicast group address
In Port	Indicates the physical interface or a logical interface (VLAN) that received group reports from various sources.

Displaying the VLAN unknown multicast no-flood status using ACLI

Use this procedure to display the status of the VLAN unknown multicast no-flood parameter. If no-flood is enabled, MAC or IP addresses cannot be configured to flood VLANs with unknown multicast packets. If no-flood is disabled, MAC or IP addresses can be configured to flood VLANs with unknown multicast packets.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show vlan igmp unknown-mcast-no-flood
```

Flushing the IGMP router table using ACLI

Use this procedure to flush members from the IGMP router table.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
ip igmp flush vlan <vid> {grp-member | mrouter}
```

Variable definitions

Variable	Value
grp-member	Flushes the IGMP group member.
mrouter	Flushes the IGMP mrouter member.

Deleting an IGMP filter profile from an Ethernet interface using ACLI

Use this procedure to remove an IGMP filter profile from a specific Ethernet interface or all Ethernet interfaces.

Prerequisites

- Log on to the Ethernet Interface Configuration mode in ACLI.

Procedure steps

At the Ethernet Interface Configuration prompt, enter either of the following commands:

```
no ip igmp filter <1-65535>
```

OR

```
default ip igmp filter <1-65535>
```

Variable definitions

Variable	Value
<1-65535>	Specifies an IGMP filter profile ID. Values range from 1 to 65535.

SSM configuration using ACLI

Displaying the SSM for IGMP configuration using ACLI

Use this procedure to display current SSM for IGMP configuration information.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp ssm
```

Job aid: command output

The following example shows sample screen output for the **show ip igmp ssm** command.

```
ERS-4524GT#show ip igmp ssm
SSM Dynamic Learning:  Enabled
SSM Range Group/Mask:  232.0.0.0/255.0.0.0
```

Configuring SSM for IGMP using ACLI

Use this procedure to configure Source-Specific Multicast (SSM) for IGMP.

Prerequisites

- Use this command in the Global Configuration mode.

Procedure steps

```
ip igmp ssm [dynamic-learning] {group-range <A.B.C.D/0-32>}
```

Variable definitions

Variable	Value
dynamic-learning	Enables or disables SSM globally.

Variable	Value
	When enabled, the switch can learn the multicast source dynamically from the IGMP proxy report.
group-range <A.B.C.D/0–32>	Specifies the IP multicast group address range source IP address and subnet mask. • A.B.C.D—IP address • 0–32—subnet mask

Disabling SSM for IGMP using ACLI

Use this procedure to disable SSM functionality for IGMP.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter the following command:

```
no ip igmp ssm [dynamic-learning]
```

Variable definitions

Variable	Value
dynamic-learning	Disables the ability for the switch to dynamically learn the multicast source IP address from the IGMP proxy report.

Setting SSM for IGMP to default using ACLI

Use this procedure to set SSM functionality for IGMP to default values.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter the following command:

```
default ip igmp ssm [dynamic-learning] [group-range]
```

Variable definitions

Variable	Value
dynamic-learning	Sets SSM dynamic learning to default (enabled). When enabled, the switch can learn the multicast source dynamically from the IGMP proxy report.
group-range	Sets the IP multicast group address range source IP address and subnet mask to default values (232.0.0.0/255.0.0.0).

SSM map configuration using ACLI

Displaying the SSM mapping table using ACLI

Use this procedure to display the SSM map configuration status and activity for IGMP.

Prerequisites

- Log on to the Privileged EXEC mode in ACLI.

Procedure steps

At the Privileged EXEC prompt, enter the following command:

```
show ip igmp ssm-map
```

Job aid: command output

The following describes SSM mapping table column headings displayed using the **show ip igmp ssm-map** command.

Heading	Description
IpMulticastGrp	Lists multicast group IP addresses.
IpSource	Lists SSM map source IP addresses.

Heading	Description
LearningMode	Indicates whether SSM traffic is statically or dynamically forwarded to IP multicast groups.
Activity	Displays SSM map activity.
AdminState	Indicates whether an SSM map is enabled or disabled.

Configuring an SSM map for IGMP using ACLI

Use this procedure to create or modify an SSM map for IGMP.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter the following command:

```
ip igmp ssm—map {<MulticastGrp_IpAddr> | <Source_IpAddr> | all}
enable
```

Variable definitions

Variable	Value
<MulticastGrp_IpAddr>	Specifies the multicast group IP address.
<Source_IpAddr>	Specifies the SSM map source IP address.
all	Specifies all multicast groups.
enable	Enables the SSM map.

Disabling IGMP SSM mapping using ACLI

Use this procedure to disable IGMP SSM mapping for a specific IP multicast group or for all IP multicast groups in the SSM group range.

Prerequisites

- Log on to the Global Configuration mode in ACLI.

Procedure steps

At the Global Configuration prompt, enter either of the following commands:

```
no ip igmp ssm—map {<MulticastGrp_IpAddr> | all} enable
```

OR

```
default ip igmp ssm—map {<MulticastGrp_IpAddr> | all} enable
```

Variable definitions

Variable	Value
<MulticastGrp_IpAddr>	Disables IGMP SSM mapping for the specified multicast group IP address.
all	Disables IGMP SSM mapping for all IP multicast groups in an SSM group range.

Job aid: Roadmap of IGMP ACLI commands

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

Command	Parameter
VLAN Interface Configuration Mode	
`ip igmp`	last-member-query-interval <last-mbr-query-int>
	mrouter <portlist>
	proxy
	query-interval <query-int>
	Query-max-response <query-max-resp>
	robust-value <robust-val>
	router-alert
	snooping
	version <1-3>

Command	Parameter
VLAN Interface Configuration Mode	
Global Configuration mode	
`vlan igmp <vid>`	unknown-mcast-no-flood {enable|disable}
	unknown-mcast-allow-flood <H.H.H> <mcast_ip_address>
Privileged EXEC mode	
`show ip igmp`	cache
	group [count][group <A.B.C.C>][member-subnet <A.B.C.D> | <0-32>]
	group-ext [count][group <A.B.C.C>] [member-subnet <A.B.C.D> | <0-32>]
	interface [vlan <vid>]
	profile <1-65535>
	router-alert [vlan <vid>]
	snooping
	ssm
	ssm-map
	unknown-mcast-allow-flood
`show vlan igmp`	unknown-mcast-allow-flood
	unknown-mcast-no-flood
`show vlan multicast membership`	<vid>
`ip igmp`	flush vlan <vid> <grp-member | mrouter | sender>

To associate IGMP with an interface, in the VLAN Interface Configuration mode, enter the command **`ip igmp`** on a VLAN. To associate IGMP with an interface and configure, for example, IGMP snooping, in the VLAN Interface Configuration mode enter the command **`ip igmp snooping`**.

To dissociate IGMP from one interface, from the VLAN Interface Configuration mode enter the command **`no ip igmp`**. This command disables snooping or proxy if they are enabled, and will delete any saved parameters for that interface.

To display only those interfaces with IGMP associated, from the VLAN Interface Configuration mode enter the command **`show ip igmp snooping/interface`**.

There is no difference in behaviour from previous releases when you use the **no ip igmp** command with parameters, for example, snooping, or when you use the **default ip igmp** command.

How to enable IGMP snooping on VLAN 2, associate IGMP with VLAN 1, and dissociate IGMP from both interfaces

From the VLAN Interface Configuration mode, enter the command **show vlan**.

The following table demonstrates the output of the **show vlan** command.

ID	Name	Type	Protocol	PID	Active	IVL/SVL Mgmt	Mgmt	Port Members	Total VLANs
1	VLAN #1	Port	None	0x0000	Yes	IVL	Yes	1-50	2

From the VLAN Interface Configuration mode, enter the command **show ip igmp interface**. Since no VLANs are associated with IGMP, there is no output for this command.

To associate IGMP with VLAN 1, enable IGMP snooping, and automatically associate IGMP on VLAN 2, enter the following commands:

- From the Configuration prompt, enter **interface vlan 1** to access the VLAN Interface Configuration mode
- From the VLAN Interface Configuration prompt, enter **ip igmp**.
- Enter the **exit** command to return to the Configuration mode..
- From the Configuration prompt, enter **interface vlan 2** to access the VLAN Interface Configuration mode.
- From the VLAN Interface Configuration prompt, enter **ip igmp snooping**.

To display information about both interfaces, from the Configuration mode, enter the command **show ip igmp interface**.

The following table demonstrates the output of the **show ip igmp interface** command.

Send VLAN Query	Query Intvl	Vers	Oper Vers	Querier	Query MaxRspT	Wrong Query	Joins	Robust	Query	Last Mbr
1	125	2	2	0.0.0.0	100	0	0	2	10	No
2	125	2	2	0.0.0.0	100	0	0	2	10	No

To dissociate IGMP from the 2 interfaces, do the following:

- From the Configuration prompt, enter **interface vlan 1**.
- From the VLAN Interface Configuration prompt, enter **no ip igmp**.

- Enter the **exit** command.
- From the Configuration prompt, enter **interface vlan 2**.
- From the VLAN Interface Configuration prompt, enter **no ip igmp**.
- From the VLAN Interface Configuration prompt, enter **show ip igmp interface**.

Because the interfaces are not associated with IGMP, the output of the command is blank.

Chapter 20: IP routing configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure routable VLANs using Enterprise Device Manager (EDM).

The Avaya Ethernet Routing Switch 4000 Series are Layer 3 switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address is attached to the VLAN. When routing is enabled in Layer 3 mode, every Layer 3 VLAN is capable of routing as well as carrying the management traffic. You can use any Layer 3 VLAN instead of the Management VLAN to manage the switch.

Configuring routing globally using EDM

Use the following procedure to configure routing at the switch level. By default, routing is disabled.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **IP**.
3. In the work area, click the **Globals** tab.
4. On the Globals tab, select the **forwarding** option in the **Forwarding** section to enable global routing.
5. Type the value of ARP lifetime in the **ARPLIfeTime** box.
6. Type the value of maximum number of local routes in the **Local** box of the **Maximum number of routes** section.
7. Type the value of maximum number of static routes in the **Static** box of the **Maximum number of routes** section.
8. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to help you configure IP routing at the switch level.

Variable	Value
Forwarding	Indicates whether routing is enabled (forwarding) or disabled (not-forwarding) on the switch.
DefaultTTL	Indicates the default time-to-live (TTL) value for a routed packet. TTL is the maximum number of seconds elapsed before a packet is discarded. The value is inserted in the TTL field of the IP header of datagrams when one is not supplied by the transport layer protocol. The TTL field is also reduced by one each time the packet passes through a router. Range is 1–255. Default value is 64 seconds.
ReasmTimeout	Indicates the maximum number of seconds that received fragments are held while they await reassembly at this entity. Default value is 60 seconds.
ARPLifeTime	Specifies the lifetime (in minutes) of an ARP entry within the system. Range is 5–360. Default value is 360 minutes.
Maximum number of routes	Specifies the maximum number of route entries. • Local—specifies the maximum number of local routes. Values range from 2–256. • Static—specifies the maximum number of static routes. Values range from 0–256. • Total—indicates the maximum sum of local, static, and dynamic routes allowed. **Important:** If you increase the combined maximum totals of local and static routes, the total number of available dynamic routes is decreased. You can decrease the total number of available dynamic routes only when IP routing is disabled. For more information, see *Dynamic Routing Table Allocation* in the *Configuring IP Routing and Multicast on Avaya Ethernet Routing Switch 4000 Series*, NN47205-506
Current allocated routes	Specifies the number of route entries currently allocated for the switch. • Local—specifies the number of local routes allocated for the switch. • Static—specifies the number of static routes allocated for the switch. • Total—specifies the combined number of local and static routes allocated for the switch.

Viewing VLAN IP Addresses using EDM

Use the following procedure to display IP address information for VLANs configured on the switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the work area, click the **Addresses** tab to display IP address information for VLANs configured on the switch.

Variable definitions

The following table describes the Addresses tab fields.

Field	Description
IfIndex	Specifies the VLAN name.
IpAddress	Specifies the associated IP address.
NetMask	Specifies the subnet mask.
BcastAddrFormat	Specifies the format of the IP broadcast address.
ReasmMaxSize	Specifies the size of the largest IP datagram that this entity can reassemble from fragmented datagrams received on this interface.
VlanId	Specifies the VLAN ID. A value of –1 indicates that the VLAN ID is ignored.
MacOffset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address.

Displaying IP routes using EDM

Use the following procedure to display information about the routes configured on your switch.

Important:

Use the following procedure to display information about the routes configured on your switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the work area, select the **Routes** tab to display the information for the routes configured on the switch.

Variable definitions

The following table describes the fields of the **Routes** tab.

Variable definition

Variable	Value
Dest	Specifies the destination address of the route.
Mask	Specifies the subnet mask for the route.
NextHop	Specifies the next hop in the route.
HopOrMetric	Specifies the metric associated with the route.
Interface	Specifies the interface associated with the route.
Proto	Specifies the protocol associated with the route. Available options are – local and static.
PathType	Specifies the route path type: • i: indirect • d: direct • A: alternative • B: best • E: ECMP • U: unresolved
Pref	Specifies the preference value associated with the route.

Configuring ECMP using EDM

Use the following procedure to configure ECMP settings for RIP, OSPF, and static routes.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Configure routing (RIP, OSPF, or static routes) on the switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the work area, click the **ECMP** tab.
4. In the table, double-click the cell under the **MaxPath** column heading for the parameter you want to change.
5. Type a numerical value from 1 to 4.
6. Repeat steps **4** and **5** as required.
7. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields on the ECMP tab.

Variable definition

Field	Description
RoutingProtocol	Indicates the routing protocol to be configured.
MaxPath	Indicates the maximum number of ECMP paths assigned to the protocol as a value in a range from 1 to 4. DEFAULT: 1

Configuring a brouter port using EDM

Use the following procedure to configure and manage brouter ports.

Procedure steps

1. In the Device Physical View, select a port.
2. Right-click the selected port.
3. Select **Edit** from the shortcut menu.

   The **Port** tab appears.
4. In the work area, click the IP Address tab.
5. In the toolbar, click **Insert**.

   The Insert IP Address dialog appears.
6. Using the provided fields, create the new brouter port.
7. Click **Insert**.

Variable definitions

Variable definition

Variable	Value
IpAddress	Specifies the IP address assigned to this brouter.
NetMask	Specifies the subnet mask associated with the brouter IP address.
VlanId	Specifies the VLAN ID associated with this brouter port.
MacOffset	Specifies the MAC address offset associated with this brouter port.

Chapter 21: Static route configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure static routes using Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.
- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLANs to be routed.

Configuring static routes using EDM

Use the following procedure to configure static routes for the switch.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **IP**.
3. In the work area, select the **Static Routes** tab.
4. On the toolbar, click **Insert**.

   The Insert Static Routes dialog box appears.
5. Type the following information for the new static route in the boxes provided.
   - **Dest**—the destination IP address.

- **Mask**—the destination mask.
- **NextHop**—the IP address of the next hop.
- **Metric**—the cost of the static route.

6. Click **Insert**.
7. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Static Routes tab fields.

Variable	Value
Dest	Specifies the destination IP address of the route. The default route is 0.0.0.0.
Mask	Specifies the destination mask of the route.
NextHop	Specifies the IP address of the next hop of this route.
Metric	Represents the cost of the static route. It is used to choose the best route (the one with the smallest cost) to a certain destination. The range is 1–65535. If this metric is not used, the value is set to –1.
IfIndex	Specifies the interface on which the static route is configured.
Enable	Specifies whether the route is administratively enabled (true) or disabled (false).
Status	Specifies the operational status of the route.

IP route information display using EDM

Use the information in this section to display general and specific IP route information.

IP route information display using EDM navigation

- Viewing IP routes using EDM on page 305
- Filtering IP route information using EDM on page 306

Viewing IP routes using EDM

Use the following procedure to display information for routes configured on the switch.

Important:

Routes are not displayed until at least one port in the VLAN has a link.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **IP**.
3. In the work area, click the **Routes** tab to display the information of the routes configured on the switch.

Variable definitions

Use the information in the following table to help you understand information displayed for routes configured on the switch.

Variable	Value
Dest	Specifies the destination address of the route.
Mask	Specifies the subnet mask for the route.
NextHop	Specifies the next hop for the route.
HopOrMetric	Specifies the metric associated with the route.
Interface	Specifies the interface associated with the route.
Proto	Specifies the protocol associated with the route. Available options are—local and static.
PathType	Specifies the route path type: • i: indirect • d: direct • B: best • U: unresolved
Pref	Specifies the preference value associated with the route.

Filtering IP route information using EDM

Use the following procedure to filter specific IP route information to display.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **IP**.
3. In the work area, click the **Routes** tab.
4. On the toolbar, click **Filter**.
5. Configure the route filter as required.
6. Click **Filter**.
7. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to help you configure route filters.

Variable	Value
Condition	Indicates the condition used to join multiple filter expressions together.
Ignore Case	Indicates whether filters are case sensitive or insensitive.
Column	Indicates the type of criteria to apply to values used for filtering.
All Records	Select this check box to clear the filters, and display all rows.
Dest	Select this check box to type a value to filter on the route destination value.
Mask	Select this check box to type a value to filter on the route destination subnet mask value.
NextHop	Select this check box to type a value to filter on the route next hop value.
HopOrMetric	Select this check box to type a value to filter on the hop count or metric of the route.
Interface	Select this check box to type a value to filter on the interface associated with the route.
Proto	Select this check box to type a value to filter on the route protocol.

Variable	Value
PathType	Select this check box to type a value to filter on the route path type.
Pref	Select this check box to type a value to filter on the route preference value.

Viewing TCP information for the switch using EDM

Use the following procedure to display Transmission Control Protocol (TCP) information for the switch.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **TCP/UDP**.
3. In the work area, click the **TCP Globals** tab to display TCP information for the switch.

Variable definitions

Use the information in the following table to understand the displayed TCP/UDP information.

Field	Description
RtoAlgorithm	Specifies the algorithm used to determine the timeout value used for retransmitting unacknowledged octets.
RtoMin	Specifies the minimum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.
RtoMax	Specifies the maximum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.
MaxConn	Specifies the limit on the total number of TCP connections that the entity can support. In entities where the maximum number of connections is dynamic, this object contains the value –1.

Viewing TCP connections using EDM

Use the following procedure to display information about the current TCP connections that the switch maintains.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **TCP/UDP**.
3. In the work area, click the **TCP Connections** tab to display information about the current TCP connections.

Variable definitions

Use the information in the following table to understand the displayed TCP connection information.

Variable	Value
LocalAddressType	Specifies the local IP address type for this TCP connection.
LocalAddress	Specifies the local IP address for this TCP connection. In case of a connection in the listen state, which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.
LocalPort	Specifies the local port number for this TCP connection.
RemAddressType	Specifies the remote IP address type for this TCP connection.
RemAddress	Specifies the remote IP address for this TCP connection.
RemPort	Specifies the remote port number for this TCP connection.
State	Specifies the state of this TCP connection.

Viewing TCP Listeners using EDM

Use the following procedure to display information about the current TCP listeners on the switch.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **TCP/UDP**.
3. In the work area, click the **TCP Listeners** tab to display information about the current TCP listeners on the switch.

Variable definitions

Use the information in the following table to understand the displayed TCP listener information.

Variable	Value
LocalAddressType	Specifies the IP address type of the local TCP listener.
LocalAddress	Specifies the local IP address of the TCP listener. The value of this field can be represented in three possible ways, depending on the characteristics of the listening application: 1. For an application willing to accept both IPv4 and IPv6 datagrams, the value of this object is a zero-length octet string, and the value of the corresponding LocalAddressType field is unknown. 2. For an application willing to accept either IPv4 or IPv6 datagrams, the value of this object must be 0.0.0.0 or ::, with the LocalAddressType identifying the supported address type. 3. For an application that is listening for data destined only to a specific IP address, the value of this object is the specific local address, with LocalAddressType identifying the supported address type.
LocalPort	Specifies the local port number for this TCP connection

Viewing UDP endpoints using EDM

Use the following procedure to display information about the UDP endpoints currently maintained by the switch.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **TCP/UDP**.
3. In the work area, click the **UDP Endpoints** tab to display information about the UDP endpoints currently maintained by the switch.
4. On the toolbar, you can click **Refresh** to update the displayed information.

Variable definitions

The following table describes the UDP Endpoints tab fields.

Variable	Value
LocalAddressType	Specifies the local address type (IPv6 or IPv4).
LocalAddress	Specifies the local IP address for this UDP listener. In the case of a UDP listener that accepts datagrams for any IP interface associated with the node, the value 0.0.0.0 is used. The value of this field can be represented in three possible ways: 1. For an application willing to accept both IPv4 and IPv6 datagrams, the value of this object is a zero-length octet string, and the value of the corresponding LocalAddressType field is unknown. 2. For an application willing to accept either IPv4 or IPv6 datagrams, the value of this object must be 0.0.0.0 or ::, with the LocalAddressType identifying the supported address type. 3. For an application that is listening for data destined only to a specific IP address, the value of this object is the address for which this node is receiving packets, with LocalAddressType identifying the supported address type.
LocalPort	Specifies the local port number for this UDP listener.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IP address for this UDP endpoint. If datagrams from all remote systems are to be accepted, this value is a zero-length octet string. Otherwise, the address of the remote system from which datagrams are to be accepted (or to which all datagrams are to be sent) is displayed with the RemoteAddressType identifying the supported address type.

Variable	Value
RemotePort	Displays the remote port number. If datagrams from all remote systems are to be accepted, this value is zero.
Instance	Distinguishes between multiple processes connected to the same UDP endpoint.
Process	Displays the ID for the UDP process.

Chapter 22: OSPF configuration using Enterprise Device Manager

This chapter describes the procedures you can use to using Enterprise Device Manager (EDM). The Open Shortest Path First (OSPF) Protocol is an Interior Protocol (IGP) that distributes routing information between belonging to a single autonomous system (AS). Intended networks, OSPF is a link-state protocol which supports the tagging of externally-derived routing information.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally.
- Assign an IP address to the VLAN that you want to enable with OSPF. Routing is automatically enabled on the VLAN when you assign an IP address to it.

Configuring OSPF globally using EDM

Use the following procedure to configure global OSPF parameters for the switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **General** tab.
4. Configure OSPF as required.
5. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to help you configure OSPF.

Variable	Value
RouterId	Specifies the unique ID of the router in the Autonomous System.
AdminStat	Specifies the administrative status (enable or disable) of OSPF on the router.
VersionNumber	Specifies the version of OSPF running on the router.
AreaBrdRtrStatus	Specifies whether this router is an Area Border Router.
ASBrdRtrStatus	Specifies whether this router is an Autonomous System Border Router.
ExternLsaCount	Specifies the number of external (link state type 5) link-state advertisements in the link state database.
ExternLsaCksumSum	Specifies the sum of the link state checksum of the external link state advertisements contained in the link state database. This sum is used to determine if the link state database of the router changes and to compare the link state databases of two routers.
OriginateNewLsas	Specifies the number of new link state advertisements that have been originated. This number is increased each time the router originates a new link state advertisement.
RxNewLsas	Specifies the number of link state advertisements received determined to be new instantiations. This number does not include newer instantiations of self-originated link state advertisements.
10MbpsPortDefault Metric	Specifies the default metric of a 10 Mbps port. The range is 1–65535. Default value is 100.
100MbpsPortDefault Metric	Specifies the default metric of a 100 Mbps port. The range is 1–65535. Default value is 10.
1GbpsPort DefaultMetric	Specifies the default metric of a 1 Gbps port. The range is 1–65535. Default value is 1.
10GbpsPort DefaultMetric	Specifies the default metric of a 10 Gbps port. The range is 1–65535. Default value is 1.
TrapEnable	Enables or disables OSPF traps. By default, OSPF trap is disabled.

Variable	Value
AutoVirtLinkEnable	Enables or disables OSPF automatic Virtual Link creation. The default setting is disabled.
SpfHoldDownTime	Specifies the SPF Hold Down Timer value, which is an integer between 3–60. Default value is 10. The SPF runs, at most, once per hold down timer value.
OspfAction	Specifies an immediate OSPF action to take. Select runSpf, and click Apply to initiate an immediate SPF run.
Rfc1583Compatibility	Controls the preference rules used when choosing among multiple Autonomous System external link state advertisements advertising the same destination. If enable, the preference rule will be the same as specified by RFC 1583. If disable, the new preference rule, as described in RFC 2328, will be applicable. This potentially prevents the routing loops when Autonomous System external link state advertisements for the same destination have been originated from different areas.
LastSpfRun	Specifies the time when the last SPF calculation was done.

Configuring an OSPF area using EDM

Use the following procedure to configure an OSPF area.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Areas** tab.
4. On the toolbar, click **Insert**.
5. Type the unique ID for the area in the **AreaId** field.
6. Choose the area type in **ImportAsExtern** section.
7. Click **Insert**.

Variable definitions

Use the information in the following table to help you configure the fields of Areas tab.

Variable	Value
AreaId	Specifies the unique identifier for the area. Area ID 0.0.0.0 is used for the OSPF backbone.
ImportAsExtern	Specifies the area type by defining its support for importing Autonomous System external link state advertisements. The options available are: • importExternal—specifies a normal area • importNoExternal—specifies a stub area • importNssa—specifies an NSSA
SpfRuns	Specifies the number of times that the OSPF intra-area route table has been calculated using this area link state database.
AreaBdrRtrCount	Specifies the total number of Area Border Routers reachable within this area. This is initially zero, and is calculated in each SPF pass.
AsBdrRtrCount	Specifies the total number of Autonomous System Border Routers reachable within this area. This is initially zero, and is calculated in each SPF pass.
AreaLsaCount	Specifies the total number of link state advertisements in the link state database of the area, excluding Autonomous System external link state advertisements.
AreaLsaCksumSum	Specifies the sum of the link state advertisements checksums contained in the link state database of this area. This sum excludes external (link state type 5) link state advertisements. The sum can be used to determine if there has been a change in link state database of a router, and to compare the link state database of two routers.
AreaSummary	Controls the import of summary link state advertisements on an ABR into a stub area. It has no effect on other areas. If the value is noAreaSummary, the ABR neither originates nor propagates summary link state advertisements into the stub area (creating a totally stubby area). If the value is sendAreaSummary, the ABR both summarizes and propagates summary link state advertisements.

Configuring an area aggregate range using EDM

Use the following procedure to configure OSPF area aggregate ranges to reduce the number of link state advertisements that are required within the area. You can also control advertisements.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Area Aggregate** tab.
4. On the toolbar, click **Insert**.
5. Click the AreaId ellipsis (**...**), and select an AreaId.
6. Choose the type of area aggregate in **LsdbType** section.
7. Type the IP address of the network or subnetwork indicated by the aggregate range in **IpAddress** field.
8. Type the subnet mask address in **Mask** field.
9. Choose the aggregate effect in **Effect** field.
10. Type the advertisement metric associated with the aggregate in **AdvertiseMetric** field.
11. Click **Insert**.
12. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Area Aggregate tab fields.

Variable	Value
AreaID	Specifies the unique identifier of the Area this address aggregate is found in.
LsdbType	Specifies the type of address aggregate. This field specifies the link state database type that this address aggregate applies to. the available options are—summaryLink and nssaExternalLink.

Variable	Value
IpAddress	Specifies the IP address of the network or subnetwork indicated by the aggregate range.
Mask	Specifies the subnet mask that pertains to the network or subnetwork.
Effect	Specifies the aggregates effect. Subnets subsumed by aggregate ranges either trigger the advertisement of the indicated aggregate (advertiseMatching value) or result in the subnet not being advertised at all outside the area. Select one of the following types: • AdvertiseMatching: advertises the aggregate summary LSA with the same LSID • DoNotAdvertiseMatching: suppresses all networks that fall within the entire range • AdvertiseDoNotAggregate: advertises individual networks
AdvertiseMetric	Specifies the advertisement metric associated with this aggregate. Enter an integer value between 0–65535 which represents the Metric cost value for the OSPF area range.

Configuring OSPF stub area metrics using EDM

Use the following procedure to display the set of metrics that are advertised by a default area border router into a stub area to determine if you wish to accept the current values or configure new ones.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Stub Area Metrics** tab.
4. Configure the stub area metrics as required.
5. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Stub Area Metrics tab fields.

Variable	Value
AreaId	Specifies the unique ID of the stub area.
TOS	Specifies the Type of Service associated with the metric.
Metric	Specifies the metric value applied to the indicated type of service. By default, this value equals the least metric at the type of service among the interfaces to other areas.
Status	Displays the status of the entry (Active or Not Active). This field is read-only.

Configuring OSPF interfaces using EDM

Use the following procedure to configure OSPF interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Interface** tab.
4. In the table, double-click the cell below the column header you want to edit.
5. Select a parameter or value from the drop-down list.
6. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Interfaces tab fields.

Variable	Value
IpAddress	Specifies the IP address of the OSPF interface.
AreaId	Specifies the unique ID of the area to which the interface connects. Area ID 0.0.0.0 indicates the OSPF backbone.
AdminStat	Specifies the administrative status of the OSPF interface.
State	Specifies the DR state of the OSPF interface: up–DR, BDR, OtherDR; down–down, and waiting.

Variable	Value
RtrPriority	In multi-access networks, specifies the priority of the interface in the designated router election algorithm. The interface with the highest priority number is the designated router. The interface with the second-highest priority becomes the backup designated router. The value 0 signifies that the router is not eligible to become the designated router on this network. This is an integer value between 0–255. In the event of a tie in the priority value, routers use their Router ID as a tie breaker. The default value is 1.
DesignatedRouter	Specifies the IP address of the Designated Router.
BackupDesignatedRouter	Specifies the IP address of the Backup Designated Router.
Type	Specifies the OSPF interface type. The options available are—broadcast and passive.
AuthType	Specifies the interface authentication type. The options available are: none, simplePassword, or md5.
AuthKey	Specifies the interface authentication key. This key is used when AuthType is simplePassword.
PrimaryMd5Key	Specifies the MD5 primary key if it exists. Otherwise this field displays 0. This key is used when AuthType is md5.
TransitDelay	Specifies the estimated number of seconds it takes to transmit a link state update packet over this interface. This is an integer value between 0–3600.
RetransInterval	Specifies the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting database description and link state request packets. This is an integer value between 0 –3600.
HelloInterval	Specifies the interval in seconds between the Hello packets sent by the router on this interface. This value must be the same for all routers attached to a common network. This is an integer value between 1–65535.
RtrDeadInterval	Specifies the number of seconds that a neighbor waits for a Hello packet from this interface before the router neighbors declare it down. This value must be some multiple of the Hello interval and must be the same for all routers attached to the common network. This is an integer value between 0–2147483647.
PollInterval	Specifies the poll interval.
AdvertiseWhenDown	Enables (true) or disables (false) the advertisement of the OSPF interface. When enabled, even if the port or VLAN for the routing interface subsequently goes down, the switch continues to advertise the route.

Variable	Value
	Note: If a port or VLAN is not operational for the routing interface, no advertisement occurs, even if you enable the *advertise-when-down* parameter.
MtuIgnore	Specifies whether the MTU value is ignored on this interface.
Events	Specifies the number of times this OSPF interface has changed its state, or an error has occurred.

Configuring OSPF interface metrics using EDM

Use the following procedure to configure OSPF interface metrics.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **If Metrics** tab.
4. In the table, select the row you want to edit.
5. In the row, double-click the cell in the **Value** column to edit the advertised value.
6. On the toolbar, click **Apply**.

Variable definitions

The following table describes the If Metrics tab fields.

Variable	Value
IpAddress	Specifies the IP address of the interface.
TOS	Specifies the Type of Service associated with the metric.
Value	Specifies the value advertised to other areas indicating the distance from the OSPF router to any network in the range. This is an integer value between 0–65535.
Status	Displays the status of the entry (Active or not Active). This field is read-only.

Defining MD5 keys for OSPF interfaces

Use the following procedure to configure OSPF MD5 keys for OSPF interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Message Digest** tab.
4. On the toolbar, Click **Insert**.
5. Click the IpAddress ellipsis (**...**), and select an IP address.
6. Type an index value for the digest entry in **Index** field.
7. Choose the digest type in **Type** field.
8. Type a key value for the digest entry in **Key** field.
9. Click **Insert**.

Variable definitions

The following table describes the Message Digest tab fields.

Variable	Value
IpAddress	Specifies the IP address of the OSPF interface associated with the digest entry.
Index	Specifies an index value for the digest entry. This is an integer value between 1–255.
Type	Specifies the type of digest entry. Only MD5 is supported.
Key	Specifies the key value associated with the digest entry.

Displaying OSPF neighbor information

Use the following procedure to display OSPF neighbors.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Neighbors** tab.
4. Click **Refresh** to update the displayed information.

Variable definitions

The following table describes the Neighbor tab fields.

Variable	Value
IpAddr	Specifies the IP address this neighbor is using as an IP source address. On addressless links, this will not be represented as 0.0.0.0 but as the address of another of the neighbor interfaces.
AddressLessIndex	Specifies the corresponding value of the interface index on addressless links. This value is zero for interfaces having an IP address.
RouterId	Specifies the unique ID of the neighboring router in the Autonomous System.
Options	Specifies a value corresponding to the neighbor Options field.
Priority	Specifies the priority of the neighbor in the designated router election algorithm. A value of 0 indicates that the neighbor is not eligible to become the designated router on this particular network. This is a value between 0–255.
State	Specifies the state of the relationship with this neighbor.
Events	Specifies the number of times this neighbor relationship has changed state or an error has occurred.
RetransQLen	Specifies the current length of the retransmission queue.
NbmaNbrPermanence	Specifies the status of the entry. The values dynamic and permanent refer to how the neighbor came to be known.
HelloSuppressed	Specifies whether Hello packets are being suppressed to the neighbor.
InterfaceAddr	Specifies the interface address of neighbor.

Configuring an OSPF virtual link using EDM

Use the following procedure to create an OSPF virtual link.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Virtual If** tab.
4. On the toolbar, click **Insert**.
5. Type the unique ID for the area in **AreaId** field.
6. Type the router ID of the virtual neighbor in the **Neighbor** field.
7. Type the estimated transit delay time in the **Transit Delay** field.
8. Type the retransmission interval time in the **RetransInterval** field.
9. Type the time interval required to send Hello packets in **HelloInterval** field.
10. Type the waiting time of the neighbor router to receive transmitted hello packets in the **RtrDeadInterval** field.
11. Click a radio button in the **AuthType** section.
12. Click **Insert**.

Variable definitions

The following table describes the Virtual If tab fields.

Variable	Value
AreaId	Specifies the unique ID of the area connected to the interface. An area ID of 0.0.0.0 indicates the OSPF backbone.
Neighbor	Specifies the router ID of the virtual neighbor.
TransitDelay	Specifies the estimated number of seconds required to transmit a link state update packet over the virtual interface. The transit delay is expressed as an integer between 1–3600. The default value is 1.
RetransInterval	Specifies the number of seconds between link state advertisement retransmissions for adjacencies belonging to the virtual interface.

Variable	Value
	The retransmit interval is also used to transmit database description and link state request packets. The retransmit interval is expressed as an integer between 1–3600. The default value is 5.
HelloInterval	Specifies the interval, in seconds, between the Hello packets sent by the router on the virtual interface. This value must be the same for all routers attached to a common network. The hello interval is expressed as an integer between 1–65535. The default value is 10.
RtrDeadInterval	Specifies the number of seconds that a neighbor router waits to receive transmitted hello packets from this interface before the neighbor declares it down. The retransmit dead interval is expressed as an integer between 1–2147483647. The retransmit dead interval must be a multiple of the hello interval and must be the same for all routers attached to a common network. The default value is 60.
AuthType	Specifies the interface authentication type. The available authentication types are—none, simplePassword, and MD5.
AuthKey	Specifies the interface authentication key used with the simplePassword authentication type.
PrimaryMd5Key	Specifies the MD5 primary key. If no MD5 primary key exists, the value in this field is 0.
State	Specifies the OSPF virtual interface state.
Events	Specifies the number of times the virtual interface has changed state or the number of times an error has occurred.
Type	Specifies whether the virtual interface is broadcast or passive.

Defining MD5 keys for OSPF virtual links using EDM

Use the following procedure to configure OSPF MD5 keys for OSPF virtual interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Virtual If Message Digest** tab.
4. On the toolbar, click **Insert**.

5. Click AreaId ellipsis (**...**), and select an area ID.
6. Click Neighbor ellipsis (**...**), and select the IP address of neighbor router.
7. Type an index value in the **Index** field.
8. Choose the digest type in the **Type** field.
9. Type the key for the digest entry in the **Key** field.
10. Click **Insert**.

Variable definitions

The following table describes the Virtual If Message Digest tab fields.

Variable	Value
AreaId	Specifies the area ID of the area associated with the virtual interface.
Neighbor	Specifies the IP address of the neighbor router associated with the virtual interface.
Index	Specifies the index value of the virtual interface message digest entry. The value is an integer between 1–255.
Type	Specifies the type of digest entry. Only MD5 is supported.
Key	Specifies the key value associated with the digest entry.

Displaying virtual neighbor information using EDM

Use this procedure to view OSPF Virtual Neighbors information.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Virtual Neighbors** tab.
4. On the toolbar, click **Refresh** to refresh the displayed information.

Variable definitions

The following table describes the Virtual Neighbors tab fields.

Variable	Value
Area	Specifies the subnetwork in which the virtual neighbor resides.
RouterId	Specifies the 32-bit integer uniquely identifying the neighboring router in the autonomous system.
IpAddr	Specifies the IP address of the virtual neighboring router.
Options	Specifies a bit mask corresponding to the option field of the neighbor.
State	Specifies the state of the virtual neighbor relationship.
Events	Specifies the number of state changes or error events that have occurred between the OSPF router and the neighbor router.
RetransQLen	Specifies the current length of the retransmission queue (the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor).
HelloSuppressed	Specifies whether Hello packets to the virtual neighbor are suppressed or not.

Configuring OSPF host routes using EDM

Use the following procedure to create OSPF hosts routes to specify which hosts are directly attached to the router and the metrics that must be advertised for them.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Hosts** tab.
4. On the toolbar, click **Insert**.
5. Type the host IP address in the **IpAddress** field.
6. Type the configured cost of the host in the **Metric** field.
7. Click **Insert**.

Variable definitions

The following table describes the Hosts tab fields.

Variable	Value
IpAddress	Specifies the host IP address.
TOS	Specifies the configured route type of service. The value in this field should be 0 as TOS-based routing is not supported.
Metric	Specifies the configured cost of the host.
AreaID	Specifies the ID of the area connected to the host.

Displaying link state database information using EDM

Use the following procedure to display OSPF link states.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Link State Database** tab.
4. Click **Refresh** to update the displayed information.

Variable definitions

The following table describes the Link State Database tab fields.

Variable	Value
AreaId	Specifies the unique identifier of the Area from which the link state advertisement was received.
Type	Specifies the type of link state advertisement. Each link state type has a separate advertisement format.
Lsid	Specifies the Link State ID, a link state type-specific field containing either a Router ID or an IP address. This field identifies the section

Variable	Value
	of the routing domain that is being described by the advertisement.
RouterId	Specifies the unique identifier of the originating router in the Autonomous System.
Sequence	This field is used to detect old or duplicate link state advertisements by assigning an incremental number to duplicate advertisements. The higher the sequence number, the more recent the advertisement.
Age	Specifies the age of the link state advertisement in seconds.
Checksum	Specifies the checksum of the complete content of the advertisement, excluding the Age field. This field is excluded so that the advertisement's age can be increased without updating the checksum. The checksum used is the same as that used in ISO connectionless datagrams and is commonly referred to as the Fletcher checksum.

Displaying external link state database information using EDM

Use the following procedure to display the OSPF external link state database.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Ext. Link State Database** tab.
4. Click **Refresh** to update the displayed information.

Variable definitions

The following table describes the Ext. Link State Database tab fields.

Variable	Value
Type	Specifies the type of link state advertisement. Each link state type has a separate advertisement format.

Variable	Value
Lsid	Specifies the Link State ID, a link state type-specific field containing either a Router ID or an IP address. This field identifies the section of the routing domain that is being described by the advertisement.
RouterId	Specifies the unique identifier of the originating router in the Autonomous System.
Sequence	This field is used to detect old or duplicate link state advertisements by assigning an incremental number to duplicate advertisements. The higher the sequence number, the more recent the advertisement.
Age	Specifies the age of the link state advertisement in seconds.
Checksum	Specifies the checksum of the complete content of the advertisement, excluding the Age field. This field is excluded so that the advertisement's age can be increased without updating the checksum. The checksum used is the same as that used in ISO connectionless datagrams and is commonly referred to as the Fletcher checksum.
Advertisement	Specifies the hexadecimal representation of the entire link state advertisement including the header.

Displaying OSPF statistics using EDM

Use the following procedure to display OSPF statistics.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Stats** tab.
4. Values on the Stats tab refreshes automatically based on the value selected in the **Poll Interval** field.
5. Click **Clear Counters** to clear the counters and start over at zero.

Variable definitions

The following table describes the Stats tab fields.

Variable	Value
LsdbTblSize	Indicates the number of entries in the link state database.
TxPackets	Indicates the number of packets transmitted by OSPF.
RxPackets	Indicates the number of packets received by OSPF.
TxDropPackets	Indicates the number of packets dropped by OSPF before transmission.
RxDropPackets	Indicates the number of packets dropped before receipt by OSPF.
RxBadPackets	Indicates the number of bad packets received by OSPF.
SpfRuns	Indicates the total number of SPF calculations performed. This also includes the number of partial route table calculations.
BuffersAllocated	Indicates the total number of buffers allocated for OSPF.
BuffersFreed	Indicates the total number of buffers that are freed by OSPF.
BufferAllocFailures	Indicates the number of times that OSPF has failed to allocate buffers.
BufferFreeFailures	Indicates the number of times that OSPF has failed to free buffers.

Chapter 23: RIP configuration using Enterprise Device Manager

This chapter describes the procedures used to configure and manage the Routing Information Protocol (RIP) using Enterprise Device Manager (EDM). RIP is a distance vector protocol used to dynamically discover network routes based on information passed between routers in the network. RIP is useful in network environments where using static route administration is difficult.

Prerequisites

- Enable IP routing globally.
- Assign an IP address to the VLAN or brouter port that you want to enable with RIP.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

Configuring global RIP properties using EDM

Use the following procedure to configure global RIP parameters.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **RIP**.
3. From the work area, click the **Globals** tab.
4. Choose the operation status in the **Operation** field.
5. Type the update time interval in the **UpdateTime** field.
6. Type the hold-time time interval in the **HoldDownTime** field.
7. Type the global timeout interval in the **TimeOutInterval** field.

8. Type the the value of the default import metric applied to routes in the **DefImportMetric** field.
9. Click **Apply**.

Variable definitions

The following table describes the Globals tab fields.

Variable	Value
Operation	Enables or disables the operation of RIP on all interfaces. The default is disabled.
UpdateTime	The time interval between RIP updates on all interfaces. It is a global parameter for the box; it applies to all interfaces and cannot be set individually for each interface. The default is 30 seconds.
RouteChanges	The number of route changes made to the IP Route Database by RIP; does not include the refresh of a route age.
Queries	The number of responses sent to RIP queries from other systems.
HoldDownTime	Sets the length of time that RIP will continue to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds.
TimeOutInterval	Specifies the global timeout interval parameter. If a RIP router does not receive an update from another RIP router within the configured timeout period, it moves the routes advertised by the nonupdating router to the garbage list. The timeout interval must be greater than the update timer. Range is 15–259200 seconds. Default is 180 seconds.
DefImportMetric	Sets the value of the default import metric applied to routes imported the RIP domain. For announcing OSPF internal routes into a RIP domain, if the policy does not specify a metric value, the default import metric is used. For OSPF external routes, the external cost is used.

Configuring a RIP interface using EDM

Use the following procedure to configure a RIP interface to tailor RIP to the individual interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, double-click **RIP**.
3. In the work area, click the **Interface** tab.
4. In the table, select the IP address row.
5. In the IP address row, double-click the cell below the **Send** or **Receive** to update the sent or received RIP version.
6. Click **Apply**.

Variable definitions

The following table describes the Interface tab fields.

Variable	Value
Address	Specifies the IP address of the RIP interface. This field is for organizational purposes only and cannot be edited.
Send	Sets the RIP version sent on this interface. The following values are valid: • doNotSend—No RIP updates sent on this interface. • ripVersion1—RIP updates compliant with RFC 1058. • rip1Compatible—Broadcasts RIPv2 updates using RFC 1058 route subsumption rules. • ripVersion2—Multicasting RIPv2 updates. The default is rip1Compatible.
Receive	Sets the RIP version received on this interface. The following values are valid: • rip1 • rip2 • rip1OrRip2 The default is rip1OrRip2. The rip2 and rip1OrRip2 imply reception of multicast packets.

Configuring advanced RIP interface properties using EDM

Use the following procedure to configure advanced RIP interface properties to fine tune and further configure a RIP interface.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **RIP**.
3. In the work area, click the **Interface Advance** tab.
4. In the table, double-click the cell below the header column you want to modify.
5. Select a parameter or value from the drop-down list.
6. Click **Apply**.

Variable definitions

The following table describes the Interface Advance tab fields.

Variable	Value
Address	Specifies the IP address of the RIP interface. This field is for organizational purposes only and cannot be edited.
Interface	Specifies the switch interface that corresponds to the listed IP address.
Enable	Enables or disables RIP on this interface.
Supply	Determines whether this interface supplies RIP advertisements.
Listen	Determines whether this interface listens for RIP advertisements.
Poison	Enables or disables poison reverse on this interface.
DefaultSupply	Determines whether this interface advertises default routes.
DefaultListen	Determines whether this interface listens for default route advertisements.
TriggeredUpdate	Enables or disables triggered updates on this interface.
AutoAggregate	Enables or disables auto aggregation on this interface.

Variable	Value
InPolicy	Associates a previously configured switch policy with this interface for use as an in policy.
OutPolicy	Associates a previously configured switch policy with this interface for use as an out policy.
Cost	The cost associated with this interface.
HoldDownTime	Sets the hold down timer for this interface. This is an integer value in seconds between 0–360.
TimeoutInterval	Sets the timeout interval for this interface. This is an integer value between 15–259200.
ProxyAnnounceFlag	Enables or disables proxy announcements on this interface.

Displaying RIP statistics using EDM

Use the following procedure to display RIP statistics.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **RIP**.
3. In the work area, click the **Stats** tab.
4. In the table, select an interface row.
5. On the toolbar, click **Graph**.
6. The table data refreshes automatically based on the value selected in the **Poll Interval** field.
7. Click **Clear Counters** to clear the counters and start over at zero.

Variable definitions

Use the data in the following table to help you understand the RIP statistics display.

Variable	Value
Address	Indicates the IP address of the RIP interface.

Variable	Value
RcvBadPackets	Indicates the number of RIP response packets received by the interface that have been discarded.
RcvBadRoutes	Indicates the number of RIP routes received by the interface that have been ignored.
SentUpdates	Indicates the number of triggered RIP updates actually sent on this interface. This does not include full updates sent containing new information.

Chapter 24: VRRP configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure Virtual Router Redundancy Protocol (VRRP) using Enterprise Device Manager (EDM).

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN you want to enable with VRRP.

Routing automatically enables on a VLAN with an assigned IP address.

Assigning a virtual router IP address using EDM

Use the following procedure to associate an IP address with a virtual router ID on a switch interface.

Procedure steps

1. From the navigation tree, double click **IP**.
2. In the IP tree, click **VRRP**
3. In the work area, click the **Interface Address** tab.
4. On the toolbar, click **Insert**.
5. In the **Index** box, enter an index value.

   OR

   Click the **VLAN** button to select a previously configured interface from the list.
6. In the **Vrid** box, enter a virtual router ID for the interface.
7. In the **IpAddr** box, enter an IP address for the interface.
8. Click **Insert**.
9. On the toolbar, click **Apply**.

Variable definition

Variable	Value
Index	The interface index for the new interface.
VrId	The virtual router ID for the interface.
IpAddr	The IP address for the interface.
Status	Indicates the status of the interface, active or inactive.

Deleting a virtual router IP address using EDM

Use this procedure to remove VRRP interface addresses.

Procedure steps

1. From the navigation tree, double click **IP**.
2. In the IP tree, click **VRRP**.
3. In the work area, click the **Interface Address** tab.
4. Select the interface you want to remove.
5. On the toolbar, click **Delete**.

Configuring VRRP globally using EDM

Use the following procedure to configure VRRP globally for the switch.

Procedure steps

1. From the navigation tree, double click **IP**.
2. In the IP tree, click **VRRP**.
3. In the work area, click the **Globals** tab.
4. Select the **Enabled** check box to enable VRRP.

   OR

   Clear the **Enabled** check box to disable VRRP.
5. Select a **NotificationCntl** button to enable or disable SNMP traps.
6. Select the **PingVirtualAddrEnabled** check box to enable virtual router ping response.

   OR

Clear the **PingVirtualAddrEnabled** check box to disable virtual router ping response.

7. On the toolbar, click **Apply**.

Variable definition

Variable	Value
Enabled	Specifies if VRRP is globally enabled.
Version	Indicates the VRRP version supported.
NotificationCntl	Specifies if the VRRP router generates SNMP traps based on VRRP events. • Enabled (checked)—send SNMP traps • Disabled (unchecked)—do not send SNMP traps
PingVirtualAddrEnabled	Indicates if this switch responds to pings sent to a virtual router IP address.

Configuring VRRP interfaces using EDM

Use this procedure to configure existing VRRP interfaces.

Procedure steps

1. From the navigation tree, double click **IP**.
2. In the IP tree, click **VRRP**.
3. In the work area, click the **Interfaces** tab.
4. In the table, double-click the cell under a column heading you wish to change.
5. Select a variable parameter or value from the drop-down list.
6. Repeat steps **4** and **5** to complete your configuration.
7. In the toolbar, click **Apply**.

Variable definition

Variable	Value
Index	The interface index of the VRRP interface.
VrId	The unique virtual router identification number.
PrimaryIpAddr	An IP address selected from the set of real interface addresses. VRRP advertisements use

Variable	Value
	the primary IP address as the source of the IP packet.
VirtualMacAddr	The virtual MAC address of the virtual router.
State	The current state of the virtual router. The states are the following: • Initialize—virtual router waiting for a startup event. • Backup—virtual router is monitoring availability of master router. • Master—virtual router is forwarding packets for associated IP addresses.
AdminState	Indicates the administrative status of the virtual router.
Priority	Indicates the priority value for the virtual router master election process, between 1 and 255. The priority value for the virtual router in master state must be 255. The default priority value for virtual routers in backup state is 100.
MasterIpAddr	Indicates real (primary) IP address of the master router. This IP address is listed as the source in the VRRP advertisement last received by this virtual router.
AdvertisementInterval	Indicates the time interval in seconds between transmissions of advertisement messages. Only the master router sends VRRP advertisements. Integer value between 1 and 255, default is 1.
VirtualRouterUpTime	Indicates the amount of time this virtual router has been running. Up time does not include initialize state.
HoldDownTimer	Indicates the time interval in seconds to wait before preempting the current master router. Integer value between 0 and 21600.
HoldDownState	The holddown state of this VRRP interface.
HoldDownTimeRemaining	Indicates the time interval in seconds before the holddown timer expires.
Action	Use to trigger an action on this VRRP interface. Options available are none (no action), or preemptHoldDownTimer.

Variable	Value
CriticalIPAddrEnabled	Indicates if the user-defined critical IP address is enabled. If disabled, the default critical IP address is 0.0.0.0.
CriticalIPAddr	The IP address of the interface to cause a shutdown event.
FastAdvertisementEnabled	Indicates if the faster advertisement interval is enabled. The default value is false (disabled).
FastAdverisementInterval	The fast advertisement time interval in milliseconds between transmissions of advertisement messages. Integer value between 200 and 1000, default is 200.

Graphing VRRP interface information using EDM

Use this procedure to view and graph VRRP statistic information.

Procedure steps

1. From the navigation tree, double click **IP**.
2. In the IP tree, click **VRRP**.
3. In the work area, click the **Interfaces** tab.
4. In the table, select an interface.
5. On the toolbar, click **Graph**.

For more information, see the following variable definitions table.

Variable definition

Variable	Value
BecomeMaster	The total number of times this virtual router has transitioned to master.
AdvertiseRcvd	The total number of VRRP advertisements received by this virtual router.
AdveritsementIntervalErrors	The total number of VRRP advertisement packets received outside of the configured advertisement interval.
IpTtlErrors	The total number of VRRP packets received by the virtual router with an IP time-to-live (TTL) not equal to 255.

Variable	Value
PriorityZeroPktsRcvd	The total number of VRRP packets received by the virtual router with a priority of 0.
PriorityZeroPktsSent	The total number of VRRP packets sent by the virtual router with a priority of 0.
InvalidTypePktsRcvd	The number of VRRP packets received by the virtual router with an invalid **type** value.
AddressListErrors	The total number of packets received with an address list not matching the locally configured list for the virtual router.
AuthFailures	The total number of VRRP packets received that do not pass the authentication check.
InvalidAuthType	The total number of packets received with an unknown authentication type.
AuthTypeMismatch	The total number of packets received with Auth Type not equal to the locally configured authentication method.
PacketLengthErrors	The total number of packets received with a packet length less than the length of the VRRP header.

Viewing general VRRP statistics using EDM

Use this procedure to display general VRRP statistic information.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **VRRP**.
3. In the work area, click the **Stats** tab.
4. On the toolbar, click **Clear Counters**.
5. On the toolbar, click the **Poll Interval** drop down menu.
6. Select a poll interval value from the list.
7. On the toolbar, click **Line**, **Area**, **Bar**, or **Pie** chart to graph the counters.

Variable definition

Variable	Value
RouterChecksumErrors	The total number of VRRP packets received with an invalid VRRP checksum value.
RouterVersionErrors	The total number of VRRP packets received with an unknown or unsupported version number.
RouterVrIdErrors	The total number of VRRP packets received with an invalid virtual router ID for this virtual router.

Chapter 25: DHCP relay configuration using Enterprise Device Manager

This chapter describes the procedures you use to configure DHCP relay using Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.
- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be set as the DHCP relay agent.
- Ensure that a route (local or static) to the destination DHCP server is available on the switch.

Configuring global DHCP Relay using EDM

Use the following procedure to configure global DHCP Relay for enabling or disabling DHCP Relay parameters for the switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the work area, click the **DHCP Relay Globals** tab.
4. Select the **DhcpForwardingEnabled** check box to enable DHCP forwarding for the switch.

5. Select the **DhcpForwardingOption82Enabled** check box to enable Option 82 for DHCP Relay.
6. Type a value in the **DhcpForwardingMaxFrameLength** box.
7. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to configure global DHCP Relay.

Variable	Value
DhcpForwardingEnabled	Enables or disables DHCP forwarding for the switch.
DhcpForwardingOption82Enabled	Enables or disables Option 82 for DHCP Relay at the switch level.
DhcpForwardingMaxFrameLength	Specifies the maximum DHCP frame length in the range of 576–1536.

Configuring DHCP Relay using EDM

Use this procedure to configure DHCP Relay.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the work area, click the **DHCP Relay** tab.
4. On the toolbar, click **Insert**.
5. Type the IP address of the local VLAN to serve as the DHCP relay agent in the **AgentAddr** box.
6. Type the remote DHCP Server IP address in the **ServerAddr** box.
7. Select the **Enable** check box.
8. Select the desired DHCP relay mode in the **Mode** section.
9. Click **Insert**.
10. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to configure DHCP Relay.

Variable	Value
AgentAddr	Specifies the IP address of the local VLAN serving as the DHCP relay agent.
ServerAddr	Specifies the IP address of the remote DHCP server.
Enable	Enables (selected) or disables (cleared) DHCP relay.
Mode	Indicates whether the relay instance applies for BOOTP packets, DHCP packets, or both.

Configuring DHCP Relay with Option 82 for a VLAN using EDM

Perform the following procedure to configure DHCP Relay with Option 82 for a VLAN.

Procedure steps

1. From the navigation tree, double-click **IP** .
2. In the IP tree, click **DHCP Relay**.
3. In the work area, click the **DHCP Relay-VLAN** tab.
4. Configure the parameters as required.
5. On the toolbar, click **Apply**.

Variable definitions

The following table describes the variables associated with DHCP parameters on VLANs.

Variable	Value
Id	Specifies an ID for the entry.
MinSec	Indicates the min-sec value. The switch immediately forwards a BootP/DHCP packet

Variable	Value
	if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped.
Enable	Specifies whether DHCP relay is enabled or disabled.
Mode	Specifies the type of packets this VLAN interface forwards: BootP, DHCP, or both.
AlwaysBroadcast	Specifies whether DHCP Reply packets are broadcast to the DHCP clients on this VLAN interface.
VlanDhcpOption82Enabled	Enables or disables option 82 on the specified VLAN. • Select **true** to enable DHCP Relay with Option 82 for the VLAN. • Select **false** to disable DHCP Relay with Option 82 for the VLAN.

Assigning an Option 82 for DHCP Relay subscriber Id to a port using EDM

Use the following procedure to assign an Option 82 for DHCP Relay subscriber Id to a port for associating an alphanumeric character string with the Option 82 function for the port.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the work area, click the **DHCP Relay-port** tab.
4. In the table, double-click the cell below the **PortDhcpOption82SubscriberId** column to edit.
5. In the cell, type a subscriber Id value for the port.
6. On the toolbar, click **Apply**.

Variable definitions

Use the data in the following table to assign a DHCP Relay Option 82 subscriber Id to a port.

Variable	Value
rcPortIndex	Indicates the slot and port number.
PortDhcpOption82SubscriberId	Specifies the DHCP Option 82 subscriber Id for the port. Value is a character string between 0–64 characters.

Viewing and graphing DHCP counters on a VLAN using EDM

Use the following procedure to display and graph the current DHCP counters on a VLAN.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click the **VLANs**.
3. In the table, click the VLAN **Id** to select a VLAN to edit.
4. On the toolbar, click **IP**.
5. In the work area, click the **DHCP** tab.
6. On the toolbar, click **Graph**.
7. On the toolbar, click **Clear Counters**.
8. On the toolbar, click the **Poll Interval** drop down menu, and then select a poll interval value.
9. On the toolbar, click **Line**, **Area**, **Bar**, **Pie**, or **Dial** chart to graph the counters.

Variable definitions

Use the information in the following table to understand the displayed and graphed DHCP counter information.

Variable	Value
NumRequests	Indicates the number of DHCP requests.
NumReplies	Indicates the number of DHCP replies.

Chapter 26: UDP broadcast forwarding configuration using Enterprise Device Manager

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address. To configure UDP broadcast forwarding using Enterprise Device Manager (EDM), follow the procedures in this chapter in the order they are presented.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.
- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a UDP forwarding interface.
- Ensure that a route (local or static) to the destination address is available on the switch.

Configuring UDP protocol table entries using EDM

Use the following procedure to create UDP table entries that identify the protocols associated with specific UDP ports that you want to forward.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the work area, click the **Protocols** tab.
4. On the toolbar, click **Insert**.

5. Type the UDP port number that you want to forward in the **PortNumber** box.
6. Type the protocol name associated with the UDP port number in the **Name** box.
7. Click **Insert**.
8. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to create UDP table entries.

Variable	Value
PortNumber	Specifies the UDP port number.
Name	Specifies the protocol name associated with the UDP port.

Configuring UDP forwarding entries using EDM

Use the following procedure to configure individual UDP forwarding entries, which associate UDP forwarding ports with destination IP addresses.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the work area, click the **Forwardings** tab.
4. On the toolbar, click **Insert**.
5. Click the **DestPort** ellipsis **(...)**, and select a destination port.
6. Type the destination address in the **DestAddr** box.
7. Click **Insert**.
8. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to configure individual UDP forwarding entries.

Variable	Value
DestPort	Specifies the port on which the UDP forwarding originates (configured using the Protocols tab).
DestAddr	Specifies the destination IP address.
Id	Specifies an ID for the entry.
FwdListIdList	Indicates the UDP forward list with which this entry is associated (using the Forwarding Lists tab).

Configuring a UDP forwarding list using EDM

Use the following procedure to add the UDP port and destination forwarding entries (configured in the Forwardings tab) to UDP forwarding lists. Each UDP forwarding list can contain multiple port/destination entries.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. IN the IP tree, click **UDP Forwarding**.
3. In the work area, click the **Forwarding Lists** tab.
4. On the toolbar, click **Insert**.
5. Type the unique ID of UDP forwarding list in the **Id** box.
6. Type a unique name for the UDP forwarding list in the **Name** box.
7. Click the FwdIdList ellipsis **(...)**, and then select the desired port and destination pairs from the list.
8. Click **OK**.
9. Click **Insert**.
10. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to add the UDP port and destination forwarding entries to UDP forwarding lists.

Variable	Value
Id	Specifies the unique identifier assigned to the forwarding list.
Name	Specifies the name assigned to the forwarding list.
FwdIdList	Specifies the forwarding entry IDs associated with the port/server IP pairs created using the Forwardings tab.

Applying a UDP forwarding list to a VLAN using EDM

Use the following procedure to assign a UDP forwarding list to a VLAN, and to configure the related UDP forwarding parameters for the VLAN.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the work area, click the **Broadcast Interfaces** tab.
4. On the toolbar, click **Insert**.
5. Click the **LocalIfAddr** ellipsis **(...)**, and then select a VLAN IP address from the list.
6. Click the **UdpPortFwdListId** ellipsis (**...**), and then select the desired UDP forwarding list to apply to the VLAN.
7. Type a numerical value in the **MaxTtl** box.
8. Type a broadcast mask value in the **BroadCastMask** box.
9. Click **Insert**.
10. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to assign a UDP forwarding list to a VLAN, and to configure the related UDP forwarding parameters for the VLAN.

Variable	Value
LocalIfAddr	Specifies the IP address of the local VLAN interface.

Variable	Value
UdpPortFwdListId	Specifies the port forwarding lists associated with the interface. This ID is defined in the Forwarding Lists tab.
MaxTtl	Indicates the maximum number of hops an IP broadcast packet can take from the source device to the destination device. The value ranges between 1–16.
NumRxPkts	Specifies the total number of UDP broadcast packets received by this local interface.
NumFwdPkts	Specifies the total number of UDP broadcast packets forwarded.
NumDropPkts DestUnreach	Specifies the total number of UDP broadcast packets dropped because the destination is unreachable.
NumDropPkts UnknownPort	Specifies the total number of UDP broadcast packets dropped because the destination port or protocol specified has no matching forwarding policy.
BroadCastMask	Specifies the 32-bit mask used by the selected VLAN interface to take forwarding decisions based on the destination IP address of the incoming UDP broadcast traffic. If you do not specify a broadcast mask value, the switch uses the mask of the interface to which the forwarding list is attached.

Chapter 27: Static ARP and Proxy ARP configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure Static ARP, display ARP entries, and configure Proxy ARP using Enterprise Device Manager (EDM).

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.
- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.

Configuring static ARP entries using EDM

Use the following procedure to configure static ARP entries for the switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the work area, click the **ARP** tab.
4. On the toolbar, click **Insert**.
5. Click **Port in Vlan**, and then select the VLAN, from the list, to which you want to add the static ARP entry.

   The Interface field updates with the appropriate VLAN and port information.
6. Type the IP address for the ARP entry in the **IPAddress** box.

7. Type the MAC address for the ARP entry in the **MacAddress** box.
8. Click **Insert**.
9. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to configure static ARP entries for the switch.

Variable	Value
Interface	Specifies the VLAN and port to which the static ARP entry is being added.
MacAddress	Specifies the MAC address of the device being set as a static ARP entry.
IpAddress	Specifies the IP address of the device being set as a static ARP entry.
Type	Specifies the type of ARP entry—static, dynamic, or local.

Configuring proxy ARP using EDM

Use the following procedure to configure proxy ARP on the switch. Proxy ARP allows the switch to respond to an ARP request from a locally attached host (or end station) for a remote destination.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the work area, click the **ARP Interfaces** tab.
4. In the table, click the VLAN ID to select a VLAN to edit.
5. In the VLAN row, double-click the cell in the **DoProxy** column.
6. Select a value from the list—**enable** to enable proxy ARP for the VLAN, or **disable** to disable proxy ARP for the VLAN.
7. In the VLAN row, double-click the cell in the **DoResp** column.

8. Select a value from the list—**enable** to enable sending ARP responses for the VLAN, or **disable** to disable sending ARP responses for the VLAN.
9. Click **Apply**.

Variable definitions

Use the information in the following table to configure proxy ARP on the switch.

Variable	Value
IfIndex	Specifies a configured switch interface.
DoProxy	Enables or disables proxy ARP on the interface.
DoResp	Specifies whether the sending of ARP responses on the specified interface is enabled or disabled.

Chapter 28: ECMP configuration using Enterprise Device Manager

This chapter describes the procedure you can use to configure ECMP using Enterprise Device Manager (EDM).

With the Equal Cost Multi Path (ECMP) feature routers can determine equal cost paths to the same destination prefix. The switch can use multiple paths for traffic load sharing and, in the event of network failure, faster convergence to other active paths. When the switch maximizes load sharing among equal-cost paths, the system uses links between routers more efficiently for IP traffic transmission

Prerequisites

- Install the Advanced License
- Open one of the supported browsers
- Enter the IP address of the switch to open an EDM session
- Enable IP routing globally
- Enable IP routing and configure an IP address on the VLANs to be routed

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **ECMP**.
3. In the work area, type a value in the MaxPath box for the desired protocol.
4. On the toolbar, click **Apply**.

Variable definitions

Use the information in the following table to configure ECMP.

Variable	Value
`MaxPath`	Specifies the number of ECMP paths allowed for a protocol, within the range 1–4.

Chapter 29: Route policies configuration using Enterprise Device Manager

This chapter describes the procedure you can use to configure route policies using Enterprise Device Manager (EDM).

Route policies are an Avaya proprietary improvement on existing routing schemes. Using existing routing schemes, packets are forwarded based on routes that have been learned by the router through routing protocols such as RIP and OSPF or through the introduction of static routes. Route policies introduce the ability to forward packets based on rule sets created by the network administrator. These rule sets, or policies, are then applied to the learned or static routes.

Prerequisites

- Open one of the supported browsers.
- Enter the IP address of the switch to open an EDM session.

Creating a prefix list using EDM

Prefix lists are the base item in a routing policy. Prefix lists contain lists of IP addresses with their associated masks that support the comparison of ranges of masks.

Use the following procedure to create a new prefix list.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **Policy**.
3. In the work area, click the **Prefix List** tab.
4. On the toolbar, click **Insert**.

5. Type a unique ID for prefix list in the **Id** field.
6. Type the IP address associated with the prefix list in the **Prefix** field.
7. Type the subnet mask length associated with the prefix list in the **PrefixMaskLen** field.
8. Type the name for the prefix list in the **Name** field.
9. Type the lower bound of the mask length in the **MaskLenFrom** field.
10. Type the upper bound of the mask length in the **MaskLenUpto** field.
11. Click **Insert**.
12. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Prefix List tab fields.

Variable	Value
Id	Specifies the unique identifier of this prefix list.
Prefix	Specifies the IP address associated with this prefix list.
PrefixMaskLen	Specifies the subnet mask length associated with this prefix list.
Name	Specifies the name associated with this prefix list.
MaskLenFrom	Specifies the lower bound of the mask length. This value, when combined with the upper bound mask length (MaskLenUpto), specifies a subnet range covered by the prefix list. The default value is the mask length (PrefixMaskLen).
MaskLenUpto	Specifies the higher bound of the mask length. This value, when combined with the lower bound mask length (MaskLenFrom), specifies a subnet range covered by the prefix list.The default value is the mask length (PrefixMaskLen).

Creating a route policy using EDM

Use the following procedure to create a new route policy. Route policies are created and then applied to the switch as accept (in), announce (out), or redistribution policies.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **Policy**.
3. In the work area, click the **Route Policy** tab.
4. Click **Insert**.
5. Type a unique policy ID in the **Id** field.
6. Type a secondary index for policy in the **SequenceNumber** field.
7. Type the policy name in the **Name** field.
8. Select **Enable** check box to enable policy sequence number.
9. Choose the mode of the policy in the **Mode** field.
10. Select the protocols to be matched in the **MatchProtocol** field.
11. Click MatchNetwork ellipsis **(...)**, and select destination network.
12. Click MatchIpRouteSource ellipsis **(...)**, and select source IP address.
13. Click MatchNextHop ellipsis **(...)**, and select next hope address.
14. Click MatchInterface ellipsis **(...)**, and select interface IP address.
15. Select the route-type to be matched for OSPF routes in the **MatchRouteType** field.
16. Type the metric for match in the **MatchMetric** field.
17. Enable or disable P bit in the **NssaPbit** field.
18. Type the route preference value in the **SetRoutePreference** field.
19. Type the route metric in the **SetMetric** field.
20. Select the type of route metric in the **SetMetricType** field.
21. Click SetInjectNetList ellipsis (**...**), and select a policy.
22. Type the route mask in the **SetMask** field.
23. Click **Insert**.
24. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Route Policy tab fields.

Variable	Value
Id	Specifies an index value to uniquely identify a policy.
SequenceNumber	Specifies a secondary index value that identifies individual policies inside a larger policy group.
Name	Specifies the name associated with this policy.
Enable	Specifies whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored.
Mode	Specifies the action to be taken when this policy is selected for a specific route. Available options are: • permit—indicates that the route is allowed. • deny—indicates that the route is ignored.
MatchProtocol	If configured, matches the protocol through which the route is learned. This field is used only for RIP announce policies. Available options are—RIP, Static, Direct, OSPF, and Any.
MatchNetwork	If configured, matches the destination network against the contents of the specified prefix list.
MatchIpRouteSource	If configured, matches the source IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
MatchNextHop	If configured, matches the next hop IP address of the route against the contents of the specified prefix list. This field applies only to non-local routes.
MatchInterface	If configured, matches the IP address of the interface by 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 type of route.
MatchRouteType	Sets a specific route-type to be matched (applies only to OSPF routes). Externaltype1 and Externaltype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes.
MatchMetric	If configured, matches the metric of the incoming advertisement or existing route against the specified value (1–655535). If set to 0, this field is ignored. The default is 0.
NssaPbit	Sets or resets the P bit in specified type 7 LSA. By default the P bit is always set in case the user sets it to a disabled state for a particular route policy than all type 7. LSAs associated with that route policy will have the P bit cleared with this intact NSSA ABR will not perform translation of these LSAs to type 5. Default is enabled.
SetRoutePreference	Specifies the route preference value to be assigned to the routes which matches this policy. This applies to Accept policies only.

Variable	Value
	You can set a value from 0–255. The default value is 0. If the default is configured, the global preference value is used.
SetMetric	If configured, the switch sets the metric value for the route while announcing or redistributing. The default-import-metric 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 the default value is used.
SetMetricType	If configured, sets the metric type for the routes to be announced into the OSPF routing protocol that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.
SetInjectNetList	If configured, the switch replaces the destination network of the route that matches this policy with the contents of the specified prefix list.
SetMask	Indicates the mask to used for routes that pass the policy matching criteria.

Configuring RIP in and out policies using EDM

Use the following procedure to configure RIP accept and announce policies.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **Policy**.
3. In the work area, click the **RIP In/Out Policy** tab.
4. In the table, in the VLAN row, double-click the cells below the **InPolicy** and **OutPolicy** to configure the RIP policies.
5. Click **Apply**.

Variable definitions

The following table describes the RIP In/Out Policy tab fields.

Variable	Value
Address	Specifies the address of the RIP interface.
Interface	Specifies the associated switch interface.
InPolicy	Specifies a previously configured policy to be used as the accept policy on this interface.
OutPolicy	Specifies a previously configured policy to be used as the announce policy on this interface.

Configuring an OSPF Accept Policy using EDM

Use the following procedure to configure OSPF accept policies.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **Policy**.
3. In the work area, click the **OSPF Accept** tab.
4. On the toolbar, click **Insert**.
5. Type the IP address of the router from which you want to accept advertisements in the **AdvertisingRtr** field.
6. Enable or disable the policy in the **Enable** field.
7. Choose the metric type in the **MetricType** field.
8. Click the PolicyName ellipsis **(...)**, and select a configured policy.
9. Click **Insert**.
10. On the toolbar, click **Apply**.

Variable definitions

The following table describes the OSPF Accept tab fields.

Variable	Value
AdvertisingRtr	Represents the IP address of the router from which advertisements are to be accepted. The value 0.0.0.0 denotes that advertisements from all routers are accepted.

Variable	Value
Enable	Indicates whether the policy is enabled.
MetricType	Indicates the metric type associated with the policy. Available options are: type1, type2, and any.
PolicyName	Specifies a previously configured policy to be used as the OSPF accept policy.

Configuring OSPF redistribution parameters using EDM

Use the following procedure to configure OSPF redistribution parameters.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **OSPF**.
3. In the work area, click the **Redistribute** tab.
4. On the toolbar, click **Insert**.
5. Choose the route source protocol in the **RouteSource** field.
6. Enable or disable the redistribution entry in the **Enable** field.
7. Type the metric in the **Metric** field.
8. Choose the metric type in the **MetricType** field.
9. Allow or suppress subnetworks in the **Subnets** field.
10. Click RoutePolicy ellipsis **(...)**, and select a preconfigured route policy to be used as the redistribution policy.
11. Click **Insert**.
12. On the toolbar, click **Apply**.

Variable definitions

The following table describes the Redistribute tab fields.

Variable	Value
RouteSource	Specifies the route source protocol for redistribution (RIP, Direct or Static).
Enable	Indicates whether the redistribution entry is active.
Metric	Specifies the metric to be announced in the advertisement. This is a value between 0–65535.
MetricType	Specifies the metric type to associate with the route redistribution—type1 or type2.
Subnets	Indicates whether subnetworks need to be advertised individually. Options available are—allow and supress.
RoutePolicy	Specifies the name of preconfigured route policy to be used as the redistribution policy.

Applying an OSPF accept or redistribution policy using EDM

Use the following procedure to configure OSPF policy application.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **Policy**.
3. In the work area, click the **Applying Policy** tab.
4. Select the **OspfInFilterApply** check box to apply a preconfigured OSPF accept policy.
5. Select the **RedistributeApply** check box to apply a preconfigured OSPF redistribution policy.
6. If you are applying OSPF redistribution policies, choose the type of redistribution to apply from the available options in the **OspfApplyRedistribute** field.
7. Click **Apply**.

Variable definitions

The following table describes the Applying Policy tab fields.

Variable	Value
OspfInFilterApply	Specifies whether OSPF accept policies are enabled.
RedistributeApply	Specifies whether OSPF redistribution policies are enabled.
OspfApplyRedistribute	Specifies the type of redistribution that is applied for OSPF redistribution policies.

Chapter 30: IGMP snooping configuration using Enterprise Device Manager

This chapter provides procedures you can use to configure the Avaya Ethernet Routing Switch 4000 Series to support IP multicast traffic using Internet Group Management Protocol (IGMP) snooping.

Displaying VLAN IGMP group information using EDM

Use the following procedure to display IGMP group information for VLANs.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs**.
3. In the work area, click the **Groups** tab.

Variable Definitions

Variable	Value
IpAddress	Indicates the multicast group Address (Class D) that others want to join. A group address can be the same for many incoming ports.
IfIndex	Indicates a physical interface or a logical interface (VLAN) that receives Group reports from various sources.
Members	Indicates the IP address of the source requesting to become a member of the IGMP group.
Expiration	Indicates the time (in seconds) remaining before the received IGMP group report expires on the port. The value is reset when a new report is received on the port.

Variable	Value
InPort	Indicates the physical or logical (VLAN) switch interface that received group reports from various sources.

Enabling or disabling unknown multicast flooding using EDM

Use this procedure to enable or disable the functionality for the switch to flood all VLANs with unknown multicast addresses.

Unknown multicast flooding is enabled by default.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **Unknown Multicast Filtering** tab.
4. To enable unknown multicast flooding, clear the **UnknownMulticastNoFlood** check box.

   OR

   To disable unknown multicast flooding, select the **UnknownMulticastNoFlood** check box.
5. Click **Apply**.

Multicast MAC address flooding using EDM

Displaying multicast MAC addresses that flood VLANs using EDM

Use this procedure to display the MAC Multicast Filter Table and view information about MAC addresses specified to flood VLANs with unknown multicast packets.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **MAC Multicast Filter Table** tab.

Variable definitions

Variable	Value
AllowedAddressVlanId	Indicates a VLAN flooded with multicast packets for a particular multicast MAC address.
AllowedAddressMacAddr	Indicates the multicast MAC address for which unknown multicast packets are flooded.

Specifying multicast MAC addresses to flood VLANs using EDM

Use this procedure to specify MAC addresses to flood VLANs with unknown multicast packets.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **MAC Multicast Filter Table** tab.
4. On the toolbar, click **Insert**.
5. In the **AllowedAddressVlanId** box, type a VLAN identifier.
6. In the **AllowedAddressMacAddr** box, type a MAC address.
7. Click **Insert**.
8. On the toolbar, click **Apply**.

Variable definitions

Variable	Value
AllowedAddressVlanId	Specifies a VLAN to flood with multicast packets for a particular multicast MAC address.
AllowedAddressMacAddr	Specifies a multicast MAC destination address for which unknown multicast packets are flooded.

Preventing multicast MAC addresses from flooding VLANS using EDM

Use this procedure to prevent MAC addresses from flooding VLANs with unknown multicast packets, by removing those MAC addresses from the MAC multicast filter table.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **MAC Multicast Filter Table** tab.
4. To select a MAC address and VLAN ID to remove from the table, click a table row.
5. On the toolbar, click **Delete**.
6. Repeat steps **4** and **5** as required.

Multicast IP address flooding using EDM

Displaying multicast IP addresses that flood VLANs using EDM

Use this procedure to display the IP Multicast Filter Table and view information about IP addresses specified to flood VLANs with unknown multicast packets.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **IP Address Multicast Filter Table** tab.

Variable definitions

Variable	Value
VlanAllowedInetAddressVlanId	Indicates a VLAN flooded with unknown multicast packets destined for a particular IP address.

Variable	Value
VlanAllowedInetAddressType	Indicates the type of IP address flooding the indicated VLAN with unknown multicast packets. Values include: • unknown • ipv4—specifies an IPv4 address type. • ipv6—specifies an IPv6 address type. • ipv4z—specifies a zoned IPv4 address type. • ipv6z—specifies a zoned IPv6 address type. • dns—specifies a Domain Name System (DNS) address type.
VlanAllowedInetAddress	Indicates the multicast group IP address for which unknown multicast traffic is flooded on the indicated VLAN.

Specifying multicast IP addresses to flood VLANs using EDM

Use this procedure to specify IP addresses to flood VLANs with unknown multicast packets.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **IP Address Multicast Filter Table** tab.
4. On the toolbar, click **Insert**.
5. In the **VlanAllowedInetAddressVlanId** box, type a VLAN identifier.
6. In the **VlanAllowedInetAddressType** section, click a radio button.
7. In the **VlanAllowedInetAddress** box, type an IP address.
8. Click **Insert**.
9. On the toolbar, click **Apply**.

Variable definitions

Variable	Value
VlanAllowedInetAddressVlanId	Specifies a VLAN flooded with unknown multicast packets destined for a particular IP address.
VlanAllowedInetAddressType	Specifies an IP address type to flood the selected VLAN with unknown multicast packets. Values include: • unknown • ipv4—specifies an IPv4 address type. • ipv6—specifies an IPv6 address type. • ipv4z—specifies a zoned IPv4 address type. • ipv6z—specifies a zoned IPv6 address type. • dns—specifies a Domain Name System (DNS) address type.
VlanAllowedInetAddress	Specifies the multicast group IP address for which unknown multicast traffic is flooded on the indicated VLAN.

Preventing multicast IP addresses from flooding VLANS using EDM

Use this procedure to prevent IP addresses from flooding VLANs with unknown multicast packets, by removing those IP addresses from the IP multicast filter table.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, click the **IP Address Multicast Filter Table** tab.
4. To select an IP address, IP address type, and VLAN ID to remove from the table, click a table row.
5. On the toolbar, click **Delete**.
6. Repeat steps **4** and **5** as required.

Configuring SSM for IGMP using EDM

Use this procedure to configure Source-Specific Multicast (SSM) for IGMP.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Globals** tab.
4. Select the **DynamicLearning** check-box to enable dynamic learning for the IGMP interface.

   OR

   Clear the **DynamicLearning** check-box to disable dynamic learning for the IGMP interface.
5. Select an **AdminAction:** radio button.
6. In the **RangeGroup** box, type an IP address.
7. In the **RangeMask** box, type a subnet mask.
8. On the toolbar, click **Apply**.

Variable definitions

Variable	Value
DynamicLearning	When selected, the switch can learn the multicast source dynamically from the IGMP proxy report.
AdminAction:	Enables or disables SSM globally.
RangeGroup	Specifies the IP multicast group address range source IP address.
RangeMask	Specifies the subnet mask for the IP multicast group address range source IP address.

SSM map configuration using EDM

Displaying the SSM mapping table using EDM

Use this procedure to display the SSM map configuration status and activity for IGMP.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **SSM Map** tab.

Variable definitions

Variable	Value
IpMulticastGrp	Indicates the multicast group IP address.
IpSource	Indicates the SSM map source IP address.
LearningMode	Indicates whether SSM traffic is statically or dynamically forwarded to the IP multicast group.
Activity	Displays SSM map activity.
AdminState	Indicates whether SSM mapping is enabled or disabled.

Creating an SSM map for IGMP using EDM

Use this procedure to create an SSM map for individual IP multicast group and IP source address pairs.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **SSM Map** tab.
4. On the menu bar, click **Insert**.

5. In the **IpMulticastGrp** box, type an IP address.
6. In the **IpSource** box, type an IP address.
7. Click **Insert**.
8. On the menu bar, click **Apply**.

Variable definitions

Variable	Value
IpMulticastGrp	Specifies the multicast group IP address.
IpSource	Specifies the SSM map source IP address.
LearningMode	Indicates whether SSM traffic is statically or dynamically forwarded to the IP multicast group.
AdminState	Indicates whether SSM mapping is enabled or disabled.

Modifying an SSM map using EDM

Use this procedure to modify the configuration of an existing SSM map.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **SSM Map** tab.
4. In the row for the map you want to edit, double-click the cell in the **IpMulticastGrp** column.
5. Type an IP address for the multicast group.
6. In the row for the map you want to edit, double-click the cell in the **IpSource** column.
7. Type an IP address for tthe SSM map source.
8. On the menu bar, click **Apply**.

IGMP interface configuration using EDM

Displaying IGMP interface configuration information using EDM

Use this procedure to display the configuration status of IGMP interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**
2. In the IP tree, double-click **IGMP**.
3. In the work area, click the **Interface** tab.

Variable definitions

Variable	Value
IfIndex	Indicates the interface on which IGMP is enabled.
QueryInterval	Indicates the frequency (in seconds) at which IGMP host query packets are transmitted on the interface. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 1–65535, and the default is 125.
Status	Indicates whether or not the interface is active. The interface becomes active if any IGMP forwarding ports exist on the interface. If the VLAN has no port members or if all of the port members are disabled, the status is notInService.
Version	Indicates the version of IGMP (1, 2, or 3) configured on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	Indicates the version of IGMP currently running on this interface.
Querier	Indicates the address of the IGMP querier on the IP subnet to which this interface is attached..
QueryMaxResponseTime	Indicates the maximum response time (in 1/10 seconds) advertised in IGMP general queries on this interface.
WrongVersionQueries	Indicates the number of queries received with an IGMP version that does not match the interface. IGMP requires that all routers on a LAN be configured to run the same version of IGMP. If

Variable	Value
	queries are received with the wrong version, it indicates a version mismatch.
Joins	Indicates 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	Indicates tuning 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 you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 2 to 255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
LastMembQueryIntvl	Indicates the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This parameter 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 seconds. Avaya recommends configuring this parameter to values higher than 3. If a fast leave process is not required, Avaya recommends values above 10. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)
RouterAlertEnable	Indicates whether router alert is enabled or disabled. 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. 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.
SendQuery	Indicates whether send query is enabled or disabled.
FlushAction	Indicates the type of IGMP router table to flush. Values include: • none • flushGrpMem—group member table • flushMrouter—mrouter table

Creating an IGMP VLAN interface using EDM

Use this procedure to create a new IGMP interface.

Important:

You can create a maximum of 256 IGMP VLAN interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. On the menu bar, click **Insert**.
5. Click the **Vlan** button to the right of the **IfIndex** box.
6. Select a VLAN interface from the list.
7. Click **Ok**.
8. In the **QueryInterval** box, type a value.
9. In the **Version** section, click a radio button.
10. In the **QueryMaxResponseTime** box, type a value.
11. In the **Robustness** box, type a value.
12. In the **LastMembQueryIntvl** box, type a value.
13. Select the **SendQuery** check-box, to enable IGMP send-query.

    OR

    Clear the **SendQuery** check-box, to disable IGMP send-query.
14. Click **Insert**.
15. On the menu bar, click **Apply**.

Variable definitions

Variable	Value
IfIndex	Specifies the interface on which IGMP is enabled.
QueryInterval	Specifies the frequency (in seconds) at which IGMP host query packets are transmitted on the interface. Values range from 1 to 65535. The default value is 125.
Version	Selects the version of IGMP (1, 2, or 3) to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
QueryMaxResponseTime	Specifies the maximum response time (in 1/10 seconds) advertised with IGMP general queries on this interface.

Variable	Value
Robustness	Specifies the tuning for the expected packet loss of a network. The robustness value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). Values range from 2 to 255. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
LastMembQueryIntvl	Specifies the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This parameter 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 seconds. Avaya recommends configuring this parameter to values higher than 3. If a fast leave process is not required, Avaya recommends values above 10. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)
SendQuery	Enables or disables IGMP send-query for the interface.

Deleting an IGMP interface using EDM

Use this procedure to remove an IGMP interface.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. To select and interface, click the **IfIndex** row.
5. On the menu bar, click **Delete**.

Modifying the IGMP query interval for an interface using EDM

Use this procedure to change the current frequency setting (in seconds) at which host query packets are transmitted on an interface.

The default query interval is 125 seconds.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the IfIndex row for the interface you want to edit, double-click the cell in the **QueryInterval** column.
5. Type a numerical value ranging from 1 to 65535.
6. Click **Apply**.

Modifying the IGMP version for an interface using EDM

Use this procedure to change the current IGMP version setting for an interface.

The default value is IGMPv2.

Important:

For IGMP to function correctly, all routers in a network must use the same version.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the IfIndex row for the interface you want to edit, double-click the cell in the **Version** column.
5. Select a version from the list.
6. Click **Apply**.

Modifying the maximum IGMP query response time using EDM

Use this procedure to change the current maximum response time setting (in 1/10 seconds) that is advertised with IGMP general queries on an interface.

The default value is 100.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.

4. In the IfIndex row for the interface you want to edit, double-click the cell in the **QueryMaxResponseTime** column.
5. Type a value in the box.
6. Click **Apply**.

Modifying IGMP robustness for an interface using EDM

Use this procedure to change the current IGMP robustness setting for an interface.

The switch uses the robustness value to offset expected packet loss on a network.

The robustness value is equal to the number of expected query packet losses for each serial query interval, plus 1.

The default value of 2 means that one query for each query interval can be dropped without the querier aging out.

Note:

Avaya recommends that you ensure the robustness value is the same as the configured value on the multicast router (IGMP querier).

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the IfIndex row for the interface you want to edit, double-click the cell in the **Robustness** column.
5. Type a numerical value from 2 to 255.
6. Click **Apply**.

Modifying the IGMP last member query interval for an interface using EDM

Use this procedure to change the maximum time interval setting (in 1/10 seconds) between group specific IGMP query messages sent on an interface, to detect the loss of the last member of an IGMP group.

The default value is 10.

Note:

Avaya recommends that you configure this parameter to values higher than 3.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the IfIndex row for the interface you want to edit, double-click the cell in the **LastMembQueryIntvl** column.
5. Type a value ranging from 0 to 255 in the box.
6. Click **Apply**.

Modifying IGMP router alert status for an interface using EDM

Use this procedure to enable or disable the ability for an interface to ignore IGMP packets that do not have the router-alert flag set in the IP header.

The default value is **disable**.

Note:

To maximize your network performance, Avaya recommends that you enable or disable IGMP router alert for the version of IGMP currently in use on the interface, as follows:

- IGMPv1— Disable
- IGMPv2—Enable
- IGMPv3—Enable

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the IfIndex row for the interface you want to edit, double-click the cell in the **RouterAlertEnable** column.
5. Select a value from the list—**enable** to enable IGMP router alert for the interface, or **disable** to disable IGMP router alert for the interface.
6. Click **Apply**.

Flushing the IGMP router table for an interface using EDM

Use the following procedure to flush a specific IGMP router table type for an interface.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the IfIndex row for the interface you want to edit, double-click the cell in the **FlushAction** column.
5. Select a value from the list.
6. Click **Apply**.

Variable definitions

Variable	Value
none	Specifies that no IGMP router table is flushed. This is the default value.
flushGrpMem:	Specifies to flush a group member table.
flushMrouter:	Specifies to flush an mrouter table.

IGMP snooping configuration for interfaces using EDM

The procedures in this section provide steps for configuring IGMP for interfaces.

Displaying the IGMP snooping configuration status for interfaces using EDM

Use this procedure to display information about the IGMP snooping configuration for interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Snoop** tab.

Variable definitions

Variable	Value
IfIndex	Indicates the VLAN ID.
SnoopEnable	Indicates the IGMP snoop status: enabled (true) or disabled (false).
ProxySnoopEnable	Indicates the IGMP proxy status: enabled (true) or disabled (false).
SnoopMRouterPorts	Indicates the static mrouter ports. Such ports are directly attached to a multicast router so the multicast data and group reports are forwarded to the router.
SnoopActiveMRouterPort	Indicates all dynamic (querier port) and static mrouter ports that are active on the interface.
SnoopMRouterExpiration	Indicates the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the interface. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Enabling or disabling IGMP snooping for interfaces using EDM

Use this procedure to enable or disable IGMP snooping for one or more interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Snoop** tab.
4. In the **IfIndex** row for the interface you want to edit, double-click the cell in the **SnoopEnable** column.
5. Select a value from the list—**true** to enable IGMP snooping for the interface, or **false** to disable IGMP snooping for the interface.
6. Repeat steps **4** and **5** for other interfaces as required.
7. Click **Apply**.

Adding static mrouter ports to interfaces using EDM

IGMP snooping considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port. By default, the switch forwards incoming IGMP membership reports only to the active mrouter port.

To forward the IGMP reports to additional ports, you can configure the additional ports as static mrouter ports.

Use this procedure to add static mrouter ports to one or more interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Snoop** tab.
4. In the **IfIndex** row for the interface you want to edit, double-click the cell in the **SnoopMRouterPorts** column.
5. To add specific mrouter ports to the interface, click the port numbers.
6. To add all available mrouter ports to the interface, click **All**.
7. Click **OK**.
8. Click **Apply**.

Enabling or disabling IGMP proxy for interfaces using EDM

Use the following procedure to enable or disable the ability for an interface to forward only specific IGMP proxy reports to the upstream mrouter.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Snoop** tab.
4. In the **IfIndex** row for the interface you want to edit, double-click the cell in the **ProxySnoopEnable** column.
5. Select a value from the list—**true** to enable IGMP proxy for the interface, or **false** to disable IGMP proxy for the interface.
6. Click **Apply**.

Displaying interface IGMP group information using EDM

Use the following procedure to display IGMP group information for interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Groups** tab

Variable definitions

Field	Description
IpAddress	Indicates the multicast group IP address.
IfIndex	Indicates the VLAN interface from which the multicast group address is heard.
Members	Indicates the IP address of the IGMP receiver (host or IGMP reporter).
Expiration	Indicates the time left before the group report expires on this port. This variable is updated upon receiving a group report.
InPort	Indicates the member port for the group. This is the port on which group traffic is forwarded.

Displaying extended interface IGMP group information using EDM

Use this procedure to display extended IGMP group information for interfaces.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Groups-Ext** tab.

Variable definitions

Variable	Value
IpAddress	Indicates the multicast group IP address.
SourceAddress	Indicates the source IP address.
Members	Indicates the IP address of the IGMP receiver (host or IGMP reporter).
Mode	Indicates the group IGMP mode.
IfIndex	Indicates the VLAN interface from which the multicast group address is heard.
Expiration	Indicates the time left before the group report expires on this port. This variable is updated upon receiving a group report.
InPort	Indicates the member port for the group. This is the port on which group traffic is forwarded.

Displaying IGMP cache information using EDM

Use this procedure to display information about the learned multicast groups in the cache and the IGMPv1 version timers

Procedure steps

1. From the navigation tree, double-click **IP**
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Cache** tab.

Variable definitions

The following table describes the fields of the Cache tab.

Field	Description
Address	Indicates the IP multicast group address.

Field	Description
IfIndex	Indicates the VLAN interface from which the group address is heard.
LastReporter	Indicates the last IGMP host to join the group.
ExpiryTime	Indicates the amount of time (in seconds) remaining before this entry is aged out..
Version1Host Timer	Indicates the time remaining until the local router assumes that no IGMP version 1 members exist on the IP subnet attached to the interface. Upon hearing an IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local interface ignores IGMPv2 Leave messages that it receives for this group.
Type	Indicates whether the entry is learned dynamically or is added statically.

IGMP profile configuration using EDM

Displaying IGMP profile information using EDM

Use this procedure to display the configuration status of IGMP profiles.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Profile** tab.

Variable definitions

Variable	Value
ProfileId	Indicates the Profile ID. The range is from 1 to 65535.
ProfileType	Indicates the type of the profile.
ProfilePortList	Indicates the list of ports to which this profile applies.

Variable	Value
ProfileDroppedPackets	Indicates the number of packets that were matched by this profile and dropped.

Creating an IGMP profile using EDM

Create an IGMP profile to configure the IGMP selective channel block feature.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Profile** tab.
4. On the toolbar, click **Insert**.
5. In the **ProfileID** dialog box, type the ProfileID.
6. Click **Insert**.

   The Profile table is updated with the created profile.
7. Double-click the cell in the **ProfilePortList** column for the new profile.
8. Select switch ports to add to the profile.
9. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the IGMP Profile tab.

Variable	Value
ProfileId	Indicates the Profile ID. Values range from 1 to 65535.
ProfileType	Indicates the type of the profile.
ProfilePortList	Specifies the list of ports to apply to this profile.
ProfileDroppedPackets	Indicates the number of packets that were matched by this profile and dropped.

Deleting an IGMP profile using EDM

Use this procedure remove an IGMP profile from the profile table.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Profile** tab.
4. Click the row for the profile you want to remove.
5. On the toolbar, click **Delete** .
6. In the confirmation field, click **Yes**.

Adding ports to an IGMP profile using EDM

Use this procedure to add ports to an existing IGMP profile.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Profile** tab.
4. In the row for the profile you want to modify, double-click the cell in the **ProfilePortList** column.
5. To add specific ports to the profile, click the port numbers.

   OR

   To add all available ports to the profile, click **All**.
6. Click **Ok**.
7. On the toolbar, click **Apply** .

Configuring an IGMP profile range using EDM

Use this procedure to set the start and end IP addresses for an IGMP profile range.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the work area, click the **Profile** tab.
4. To select a profile, click the profile row.
5. On the toolbar, click **Profile Range**.

6. In the Profile Range work area, double-click the cell under in the **RangeAddressStart** column.
7. Type an IP address.
8. In the Profile Range work area, double-click the cell under in the **RangeAddressEnd** column.
9. Type an IP address.
10. In the toolbar, click **Apply**

Variable definitions

Variable	Value
ProfileId	Indicates the Profile ID. Values range from 1 to 65535.
RangeAddressStart	Specifies the IP address for the start of the IGMP profile range.
RangeAddressEnd	Specifies the IP address for the end of the IGMP profile range.

Chapter 31: IP Routing capabilities and limitations

The following table lists the capabilities and limitations of IP Routing features and protocols for the Avaya ERS 4000 Series.

Table 11: Capabilities and limitations

Feature	Maximum number supported
ARP entries (local, static & dynamic)	1792
Local ARP Entries (local IP interfaces)	256
Static ARP entries	256
Dynamic ARP entries	1280
IPv4 route entries (local, static & dynamic)	512
Static routes	32 (configurable 0-256)
Local routes	64 (configurable up to 2-256)
Dynamic routes (RIP & OSPF)	416 (configurable up to 510)
Dynamic routing interfaces (RIP & OSPF)	64
OSPF areas	4 (3 areas plus area 0)
OSPF Adjacencies	16
OSPF Virtual Links	4
L3 VLANs supported by OSPF	256
Host routes supported by OSPF	32
Areas supported by OSPF	3 non-backbone areas and area 0
Area aggregate ranges for each area supported by OSPF	8
Management routes	4
UDP Forwarding entries	128
UDP port/protocol entries	128
VLANs bound to a single UDP forwarding list	16
Ports with IP addresses in single UDP forwarding list	16

Feature	Maximum number supported
DHCP relay entries	256
DHCP relay forward paths	512
RIP routes	510
RIP Layer 3 VLANs	256
ECMP paths	4
Miscellaneous	
When adding a static ARP entry for a VLAN subnet, the IP address associated with the MAC address must be in the subnet for the VLAN. Otherwise the following error message is returned: `% Cannot modify settings` `IP address does not match with VLAN subnet.`	

Glossary

ACLI	Avaya Command Line Interface (ACLI) is a text-based, common command line interface used for device configuration and management across Avaya products.
ACLI modes	Differing command modes are available within the text-based interface, dependant on the level of user permissions determined by logon password. Each successive mode level provides access to more complex command sets, from the most restrictive—show level only, to the highest configuration levels for routing parameters, interface configuration, and security.
Address Resolution Protocol (ARP)	Maps an IP address to a physical machine address, for example, maps an IP address to an Ethernet media access control (MAC) address.
American Standard Code for Information Interchange (ASCII)	A code to represent characters in computers. ASCII uses uppercase and lowercase alphabetic letters, numeric digits, and special symbols.
area border router (ABR)	A router attached to two or more areas inside an Open Shortest Path First (OSPF) network. Area border routers play an important role in OSPF networks by condensing the amount of disseminated OSPF information.
Autonomous System (AS)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the AS, and using an EGP to route packets to other ASs.
Autonomous System Number (ASN)	A two-byte number that is used to identify a specific AS.
Automatic PVID	Automatically sets the port-based VLAN ID when you add the port to the VLAN. The PVID value is the same value as the last port-based VLAN ID associated with the port.
backup designated router (BDR)	A router that assumes the designated router (DR) role for the Open Shortest Path First (OSPF) protocol if the DR fails.
bandwidth	A measure of transmission capacity for a particular pathway, expressed in megabits per second (Mb/s).

base unit (BU) When you connect multiple switches into a stack, one unit, and only one unit, must be designated as a base unit to perform stack configuration tasks. The position of the unit select switch, on the back of the switch, determines base unit designation.

Bootstrap Protocol (BootP) A User Datagram Protocol (UDP)/Internet Protocol (IP)-based protocol that a booting host uses to configure itself dynamically and without user supervision.

brouter port A single port VLAN that can route IP packets and bridge all non-routable traffic.

designated router (DR) A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group.

Domain Name System (DNS) A system that maps and converts domain and host names to IP addresses.

Dynamic Host Configuration Protocol (DHCP) A standard Internet protocol that dynamically configures hosts on an Internet Protocol (IP) network for either IPv4 or IPv6. DHCP extends the Bootstrap Protocol (BOOTP).

Dynamic Host Configuration Protocol relay (DHCP Relay) Allows forwarding of client requests to DHCP servers residing on different IP subnets from the client.

Dynamic Host Configuration Protocol Snooping (DHCP Snooping) Prevents DHCP Spoofing attacks by ensuring client ports can only request appropriate DHCP information and are not permitted to source DHCP leases.

Enterprise Device Manager (EDM) A Web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.

equal cost multipath (ECMP) Distributes routing traffic among multiple equal-cost routes.

Extensible Authentication Protocol over LAN (EAPoL) A port-based network access control protocol. EAPoL provides security in that it prevents users from accessing network resources before they are authenticated.

Institute of Electrical and Electronics Engineers (IEEE) An international professional society that issues standards and is a member of the American National Standards Institute, the International Standards Institute, and the International Standards Organization.

internal router (IR) A router with interfaces only within a single area inside an Open Shortest Path First (OSPF) network.

Internet Control Message Protocol (ICMP) A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.

Internet Group Management Protocol (IGMP) IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.

Internet Protocol routing (IP routing) Provides a stable route or external gateway to leave an autonomous system by suing selflearning and self-healing dynamic routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF).

Internet Protocol version 4 (IPv4) The protocol used to format packets for the Internet and many enterprise networks. IPv4 provides packet routing and reassembly.

Internet Protocol version 6 (IPv6) An improved version of the IP protocol, IPv6 improves the IPv4 limitations of security and user address numbers.

last member query interval (LMQI) The time between when the last IGMP member leaves the group and the stream stops.

latency The time between when a node sends a message and receipt of the message by another node; also referred to as propagation delay.

Layer 2 Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are: Ethernet and Frame Relay.

Layer 3 Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).

Link Aggregation Provides the mechanism to create and manage trunk groups automatically using Link Aggregation Control Protocol (LACP).

Link Aggregation Control Protocol (LACP) A network handshaking protocol that provides a means to aggregate multiple links between appropriately configured devices.

link-state advertisement (LSA) Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.

link-state database (LSDB)	A database built by each OSPF router to store LSA information. The router uses the LSDB to calculate the shortest path to each destination in the Autonomous System (AS), with itself at the root of each path.
Local Area Network (LAN)	A data communications system that lies within a limited spatial area, uses a specific user group and topology, and can connect to a public switched telecommunications network (but is not one).
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
mask	A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.
maximum transmission unit (MTU)	The largest number of bytes in a packet—the maximum transmission unit of the port.
Media Access Control (MAC)	Arbitrates access to and from a shared medium.
Message Digest 5 (MD5)	A one-way hash function that creates a message digest for digital signatures.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.
Multiple Spanning Tree Protocol (MSTP)	Configures multiple instances of the Rapid Spanning Tree Protocol (RSTP) on the switch.
Network Interface Card (NIC)	A network interface device (NID) in the form of a circuit card installed in an expansion slot of a computer to provide network access.
nonbase unit (NBU)	A nonbase unit is any unit in a stack except the base unit.
NonVolatile Random Access Memory (NVRAM)	Random Access Memory that retains its contents after electrical power turns off.
not so stubby area (NSSA)	Prevents the flooding of external link-state advertisements (LSA) into the area by providing them with a default route. An NSSA is a configuration of the Open Shortest Path First (OSPF) protocol.
Open Shortest Path First (OSPF)	A link-state routing protocol used as an Interior Gateway Protocol (IGP).

Open Systems Interconnection (OSI) A suite of communication protocols, network architectures, and network management standards produced by the International Organization for Standardization (ISO). OSI-compliant systems can communicate with other OSI-compliant systems for a meaningful exchange of information.

packet loss Expressed as a percentage of packets dropped over a specified interval. Keep packet loss to a minimum to deliver effective IP telephony and IP video services.

port A physical interface that transmits and receives data.

port mirroring A feature that sends received or transmitted traffic to a second destination.

port VLAN ID Used to coordinate VLANs across multiple switches. When you create a port-based VLAN on a switch, assign a VLAN identification number (VLAN ID) and specify the ports that belong to the VLAN.

prefix A group of contiguous bits, from 0 to 32 bits in length, that defines a set of addresses.

Proxy Address Resolution Protocol (Proxy ARP) Allows the switch to respond to an Address Resolution Protocol (ARP) request from a locally attached host (or end station) for a remote destination.

quality of service (QoS) QoS features reserve resources in a congested network, allowing you to configure a higher priority to certain devices. For example, you can configure a higher priority to IP deskphones, which need a fixed bit rate, and, split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.

Rapid Spanning Tree Protocol (RSTP) Reduces the recovery time after a network breakdown. RSTP enhances switch-generated Topology Change Notification (TCN) packets to reduce network flooding.

rate limiting Rate limiting sets the percentage of traffic that is multicast, broadcast, or both, on specified ports.

request for comments (RFC) A document series published by the Internet Engineering Task Force (IETF) that describe Internet standards.

route policies An Avaya proprietary improvement, route policies can forward packets based on rule sets created by the network administrator on routes learned through routing protocols or the introduction of static routes.

Routing Information Protocol (RIP) A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.

routing policy A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.

routing switch Virtualizes the physical router interfaces to switches. A virtual router port, or interface, acts as a router port to consolidate switching and routing functions in the broadcast domain, or between broadcast domains, and enable IP routing for higher traffic volumes.

shortest path first (SPF) A class of routing protocols that use Djikstra's algorithm to compute the shortest path through a network, according to specified metrics, for efficient transmission of packet data.

spanning tree A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.

Spanning Tree Group (STG) A collection of ports in one spanning tree instance.

Spanning Tree Protocol (STP) MAC bridges use the STP to exchange information across Local Area Networks to compute the active topology of a bridged Local Area Network in accordance with the Spanning Tree Protocol algorithm.

stack Stackable Avaya Ethernet Routing Switches can be connected in a stack configuration of two or more units, up to eight units maximum. A switch stack operates and is managed as a single virtual switch.

stack IP address An IP address must be assigned to a stack so that all units can operate as a single entity.

stack unit Any switch within a stack.

Static Address Resolution Protocol (Static ARP) When you configure a Static ARP entry, both the IP address and MAC address of a device are assigned to a physical port. You can use Static ARP entries to communicate with a device that does not respond to an ARP request and to prevent an existing ARP entry from aging out.

temporary base unit (TBU) If an assigned base unit in a stack fails, the next unit in the stack automatically becomes the temporary base unit (TBU). The TBU maintains stack operations until the stack is restarted or the TBU fails. If the old base unit rejoins the stack, it does not take over from the TBU until the stack is reset.

time-to-live (TTL) The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.

Transmission Control Protocol (TCP)	Provides flow control and sequencing for transmitted data over an end-to-end connection.
Transmission Control Protocol/ Internet Protocol (TCP/IP)	Provides communication across interconnected networks, between computers with diverse hardware architectures and various operating systems—TCP/IP signifies the family of common Internet Protocols that define the Internet. Transmission Control Protcol is connection oriented and provides reliable communication and multiplexing, and IP is a connectionless protocol providing packet routing.
Trivial File Transfer Protocol (TFTP)	A protocol that governs transferring files between nodes without protection against packet loss.
trunk	A logical group of ports that behaves like a single large port.
type of service (TOS)	A field in the IPv4 header that determines the Class of Service prior to the standardization of Differentiated Services.
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/ IP host systems use UDP for application-to-application programs.
User Datagram Protocol broadcast forwarding (UDP broadcast forwarding)	Can selectively forward limited UDP broadcasts, received on an IP interface, to a configured IP address.
Virtual Local Area Network (VLAN)	A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.