



Configuration — IP Routing Protocols Avaya Ethernet Routing Switch 5000 Series

6.3
NN47200-503
07.02
October 2012

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

Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree 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 Web sites referenced within this site or documentation(s) 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 this product. 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 Web site: <http://www.avaya.com/support>. Please note that if you acquired the product from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/](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 RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. 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 RESELLER, AND 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").

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation(s) and Product(s) provided by Avaya. All content on this site, the documentation(s) and the product(s) 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. 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

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: <http://www.avaya.com/support/Copyright/>.

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

All other trademarks are the property of their respective owners.

Downloading documents

For the most current versions of documentation, see the Avaya Support Web site: <http://www.avaya.com/support>

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: <http://www.avaya.com/support>

Contents

Chapter 1: New in this release	15
Features.....	15
Circuitless IP.....	15
Content-Based Forward to Next Hop Enhancements.....	15
IGMP functionality over split multilink trunking (SMLT) network topologies.....	16
IGMP Multicast flood control.....	17
PIM-SSM.....	17
Routing IP directed broadcasts per VLAN.....	18
VRF Lite.....	18
Chapter 2: Introduction	19
ACLI command modes.....	19
Chapter 3: IP routing fundamentals	21
IP address overview.....	21
Subnet addresses.....	22
IP routing.....	23
IP routing using VLANs.....	24
Local routes.....	24
Local and non-local static routes.....	26
Default routes.....	26
Management VLAN.....	26
Multinetting.....	29
Brouter port.....	30
RIP.....	31
RIP operation.....	31
RIP metrics.....	31
Routing updates.....	32
Split horizon.....	33
Triggered updates.....	33
RIP send and receive modes.....	33
Supported RIP capabilities on the 5000 switch.....	34
RIP limitations.....	35
OSPF.....	35
Overview.....	36
Autonomous system and areas.....	36
ASBR and external router advertisements.....	38
OSPF neighbors.....	39
Designated routers.....	39
OSPF operation.....	39
OSPF router advertisements.....	40
Router types.....	40
LSA types.....	41
Area types.....	42
OSPF virtual link.....	46
OSPF host route.....	48

OSPF interfaces.....	48
OSPF packets.....	49
OSPF metrics.....	50
Automatic router ID change in a stack.....	50
OSPF security mechanisms.....	51
Equal Cost MultiPath.....	51
Route policies.....	52
Route policies in a stack.....	53
DHCP relay.....	54
Forwarding DHCP packets.....	55
Multiple DHCP servers.....	56
Differences between DHCP and BootP.....	57
IP forwarding next-hop.....	57
Limitations.....	58
Enhancements.....	59
UDP broadcast forwarding.....	60
Directed broadcasts.....	61
Routing IP directed broadcasts per VLAN.....	61
ARP.....	62
Static ARP.....	62
Proxy ARP.....	63
IP blocking.....	64
IP blocking for stacks.....	64
Virtual Router Redundancy Protocol (VRRP).....	65
VRRP operation.....	66
VRRP topology example.....	67
Critical IP address.....	68
VRRP and SMLT.....	68
VRRP fast advertisement interval.....	70
VRF Lite.....	70
Circuitless IP.....	72
Chapter 4: IP multicast fundamentals.....	75
Overview of IP multicast.....	75
Multicast groups.....	78
Multicast distribution trees.....	78
Multicast addresses.....	80
IP multicast address ranges.....	80
IP to Ethernet multicast MAC mapping.....	81
Internet Group Management Protocol.....	82
IGMPv1 operation.....	83
IGMPv2 operation.....	84
IGMPv3 operation.....	87
IGMPv3 membership report.....	88
IGMPv3 membership query.....	89
Multicast flow over Multi-Link Trunking.....	90
IGMP requests for comment.....	90
IGMP snooping.....	90

IGMPv3 snooping.....	92
IGMP proxy.....	93
Report forwarding.....	94
Static mrouter port and nonquerier.....	95
Unknown multicast packet filtering.....	95
IGMP snooping configuration rules.....	96
IGMP and stacking.....	98
Default IGMP values.....	98
IGMP snooping interworking with Windows clients.....	98
IGMP Send Query.....	99
Protocol Independent Multicast-Sparse Mode.....	100
PIM-SM concepts and terminology.....	101
PIM-SM shared trees and shortest-path trees.....	105
Source-to-RP SPT.....	108
Register suppression timeout.....	109
Receivers leaving a group.....	109
PIM assert.....	109
PIM passive interfaces.....	110
PIM-SM capabilities and limitations.....	111
Enabling or disabling routing with IGMP enabled.....	112
Nonsupported IGMP features.....	112
Default PIM-SM values.....	112
PIM-SSM overview.....	114
PIM SSM concepts and terminology.....	115
Theory of operation.....	116
Static IP routing table.....	117
IGMP functionality over split multilink trunking (SMLT) network topologies.....	118
IGMP Multicast flood control.....	119
Chapter 5: IPv6 routing fundamentals.....	121
IPv6 routing features.....	121
Host autoconfiguration.....	121
IPv6 static routes.....	123
Management route.....	125
IPv6 DHCP Relay.....	125
IPv6-in-IPv4 tunnels.....	126
Tunneling limitations.....	128
IPv6 management.....	128
Chapter 6: IP routing configuration using ACLI.....	129
Configuring global IP routing status.....	129
Displaying global IP routing status.....	130
Configuring an IP address for a VLAN.....	130
Configuring IP routing status on a VLAN.....	131
Configuring a secondary IP address for a VLAN.....	132
Displaying the IP address configuration and routing status for a VLAN.....	133
Displaying IP routes.....	134
Performing a traceroute.....	136
Entering Router Configuration mode.....	137

Accessing Loopback Interface Configuration mode.....	137
Configuring a CLIP interface.....	138
Deleting CLIP configuration parameters.....	139
Restoring CLIP to default.....	139
Displaying CLIP information.....	140
Chapter 7: IP routing configuration using EDM.....	141
Configuring global IP routing status and ARP lifetime.....	141
Configuring IP directed broadcasts per VLAN.....	142
Configuring an IP address and enabling routing for a VLAN.....	143
Displaying configured IP Addresses.....	144
Configuring a CLIP interface.....	145
Deleting a CLIP interface.....	146
Configuring a CLIP interface for OSPF.....	146
Chapter 8: VRF Lite configuration using ACLI.....	149
Configuring a VRF instance.....	149
Deleting a VRF instance.....	150
Displaying VRF instances.....	150
Associating a VLAN with a VRF instance.....	151
Removing a VLAN association with a VRF instance.....	152
Displaying VLAN and VRF instance associations.....	152
Associating a port with a VRF instance.....	153
Removing a port association with a VRF instance.....	154
Displaying switch port and VRF instance associations.....	154
Changing the router mode to a VRF instance.....	154
Chapter 9: VRF Lite configuration using EDM.....	157
Configuring a VRF instance.....	157
Creating a VRF instance.....	158
Deleting a VRF instance.....	158
Viewing brouter port and VRF associations.....	159
Associating a VLAN with a VRF instance.....	160
Associating a switch port with a VRF instance.....	160
Selecting and launching a VRF context view.....	161
Chapter 10: Brouter port configuration using ACLI.....	163
Configuring a brouter port.....	163
Displaying the brouter port configuration.....	164
Chapter 11: Brouter port configuration using EDM.....	167
Configuring a brouter port.....	167
Chapter 12: Static route configuration using ACLI.....	169
Configuring a static route.....	169
Displaying static routes.....	170
Configuring a management route.....	171
Displaying the management routes.....	172
Displaying the multicast static IP routing table using ACLI.....	173
Creating a multicast-static IP routing table entry using ACLI.....	173
Deleting a multicast-static IP routing table entry using ACLI.....	174
Enabling a route in the multicast-static IP routing table using ACLI.....	174
Disabling a route in a multicast static IP routing table using ACLI.....	174

Modifying the administrative distance of a multicast static IP route using ACLI.....	175
Chapter 13: Static route configuration using Enterprise Device Manager.....	177
Configuring static routes.....	177
Displaying IP routes.....	178
Filtering route information.....	179
Creating a multicast-static IP routing table entry.....	180
Displaying TCP information for the switch.....	181
Displaying TCP connections.....	182
Displaying TCP listeners.....	182
Displaying UDP endpoints.....	183
Chapter 14: OSPF configuration using ACLI.....	187
Configuring the OSPF hardware mode.....	187
Configuring the router ID.....	188
Configuring global OSPF status.....	189
Configuring global OSPF parameters.....	189
Displaying global OSPF parameters.....	190
Configuring OSPF area parameters.....	191
Displaying OSPF area configuration.....	192
Displaying OSPF area range information.....	193
Enabling OSPF routing on an interface.....	193
Assigning an interface to an OSPF area.....	194
Configuring the OSPF properties for an interface.....	195
Displaying OSPF interface timers.....	196
Displaying OSPF interface configurations.....	197
Displaying OSPF neighbors.....	197
Specifying a router as an ASBR.....	198
Configuring the OSPF authentication type for an interface.....	198
Defining simple authentication keys for OSPF interfaces.....	199
Defining MD5 keys for OSPF interfaces.....	199
Displaying OSPF MD5 keys.....	200
Applying an MD5 key to an OSPF interface.....	201
Displaying OSPF interface authentication configuration.....	202
Configuring a virtual link.....	202
Creating a virtual interface message digest key.....	203
Configuring automatic virtual links.....	204
Displaying OSPF virtual links.....	207
Displaying OSPF virtual neighbors.....	207
Configuring an OSPF host route.....	207
Displaying OSPF host routes.....	208
Displaying the OSPF link state database.....	209
Displaying the external link state database.....	209
Initiating an SPF run to update the OSPF LSDB.....	210
Displaying OSPF default port metrics.....	210
Displaying OSPF statistics.....	211
Displaying OSPF interface statistics.....	211
Clearing OSPF statistics counters.....	212
Chapter 15: OSPF configuration examples using ACLI.....	213

Basic OSPF configuration examples.....	213
Advanced OSPF configuration examples.....	216
Configuring ASBRs.....	221
Configuring a multi-area complex.....	229
Diagnosing neighbor state problems.....	258
Chapter 16: OSPF configuration using Enterprise Device Manager.....	261
Configuring Global OSPF properties.....	261
Configuring an automatic virtual link.....	263
Configuring an OSPF area.....	264
Configuring an area aggregate range.....	266
Configuring OSPF stub area metrics.....	267
Configuring OSPF interfaces.....	268
Configuring OSPF interface metrics.....	270
Defining MD5 keys for OSPF interfaces.....	272
Displaying OSPF neighbor information.....	273
Configuring an OSPF virtual link.....	274
Defining MD5 keys for OSPF virtual links.....	276
Displaying virtual neighbor information.....	277
Configuring OSPF host routes.....	278
Displaying link state database information.....	279
Displaying external link state database information.....	280
Displaying OSPF statistics using EDM.....	281
Displaying VLAN OSPF statistics.....	283
Chapter 17: RIP configuration using ACLI.....	285
Configuring the global RIP status.....	285
Configuring the RIP global timeout, holddown timer, and update timer.....	286
Configuring the default RIP metric value.....	287
Displaying global RIP information.....	288
Configuring the RIP status on an interface.....	289
Configuring RIP parameters for an interface.....	289
Displaying RIP interface configuration.....	291
Triggering a RIP update manually.....	293
Chapter 18: RIP configuration examples using ACLI.....	295
Chapter 19: RIP configuration using Enterprise Device Manager.....	305
Configuring RIP using EDM.....	305
Configuring Global RIP properties using EDM.....	305
Configuring a RIP interface using EDM.....	307
Configuring advanced RIP interface properties using EDM.....	308
Displaying RIP Statistics using EDM.....	309
Configuring RIP parameters for a VLAN using EDM.....	310
Chapter 20: ECMP configuration using ACLI.....	313
Configuring the number of ECMP paths allotted for RIP.....	313
Configuring the number of ECMP paths allotted for OSPF.....	314
Configuring the number of ECMP paths allotted for static routes.....	315
Configuring the number of ECMP paths allotted for BGP.....	315
Displaying ECMP path information.....	316
ECMP configuration examples.....	317

Chapter 21: ECMP configuration using Enterprise Device Manager.....	321
Configuring ECMP using EDM.....	321
Chapter 22: Route policy configuration using ACLI.....	323
Configuring prefix lists.....	323
Configuring route maps.....	324
Displaying route maps.....	327
Applying a RIP accept (in) policy.....	328
Applying a RIP announce (out) policy.....	329
Configuring an OSPF accept policy.....	330
Applying the OSPF accept policy.....	331
Displaying the OSPF accept policy.....	331
Configuring an OSPF redistribution policy.....	332
Applying the OSPF redistribution policy.....	333
Displaying the OSPF redistribution policy.....	333
Configuring IP forwarding next-hop.....	334
Displaying the IP forwarding next-hop configuration.....	336
Chapter 23: Route policy configuration using Enterprise Device Manager.....	337
Creating a prefix list.....	337
Creating a route policy.....	338
Configuring RIP in and out policies.....	341
Configuring an OSPF Accept Policy.....	342
Configuring OSPF redistribution parameters.....	343
Applying an OSPF accept or redistribution policy.....	344
Configuring the global IP forwarding next-hop status.....	345
Configuring an IP forwarding next-hop policy.....	345
Configuring an IP forwarding next-hop policy for an interface.....	347
Chapter 24: DHCP relay configuration using ACLI.....	349
Configuring global DHCP relay status.....	349
Displaying the global DHCP relay status.....	350
Specifying a local DHCP relay agent and remote DHCP server.....	350
Configuring DHCP relay maximum frame.....	352
Displaying the DHCP relay configuration.....	352
Configuring DHCP relay status and parameters on a VLAN.....	353
Configuring DHCP relay option 82 globally.....	354
Configuring DHCP relay option 82 on a port.....	354
Displaying the DHCP relay configuration for a VLAN.....	355
Displaying DHCP relay counters.....	356
Clearing DHCP relay counters for a VLAN.....	357
Chapter 25: DHCP relay configuration using Enterprise Device Manager.....	359
Configuring global DHCP Relay status and parameters.....	359
Configuring a DHCP Relay forwarding path using EDM.....	360
Configuring DHCP parameters on a VLAN using EDM.....	361
Configuring DHCP Relay option 82 on a VLAN.....	362
Configuring DHCP Relay option 82 on a port.....	363
Chapter 26: UDP broadcast forwarding configuration using ACLI.....	365
Configuring UDP protocol table entries.....	365
Displaying the UDP protocol table.....	366

Configuring a UDP forwarding list.....	367
Applying a UDP forwarding list to a VLAN.....	367
Displaying the UDP broadcast forwarding configuration.....	368
Clearing UDP broadcast counters on an interface.....	370
Chapter 27: UDP broadcast forwarding configuration using Enterprise Device Manager.....	371
Configuring UDP protocol table entries.....	372
Configuring UDP forwarding entries.....	373
Configuring a UDP forwarding list.....	374
Applying a UDP forwarding list to a VLAN.....	375
Chapter 28: Directed broadcasts configuration using ACLI.....	377
Configuring directed broadcasts.....	377
Displaying the directed broadcast configuration.....	377
Configuring IP directed broadcasts for each VLAN.....	378
Enabling IP directed broadcasts for each VLAN.....	378
Disabling IP directed broadcasts for each VLAN.....	379
Setting IP directed broadcasts for each VLAN to default.....	379
Chapter 29: Static ARP and Proxy ARP configuration using ACLI.....	381
Static ARP configuration.....	381
Configuring a static ARP entry.....	381
Displaying the ARP table.....	382
Displaying ARP entries.....	382
Configuring a global timeout for ARP entries.....	384
Clearing the ARP cache.....	384
Proxy ARP configuration.....	385
Configuring proxy ARP status.....	385
Displaying proxy ARP status on a VLAN.....	385
Chapter 30: Static ARP and Proxy ARP configuration using Enterprise Device Manager.....	387
Configuring static ARP entries.....	387
Configuring Proxy ARP.....	388
Chapter 31: IP blocking configuration using ACLI.....	391
Configuring IP blocking for a stack.....	391
Displaying IP blocking status.....	391
Chapter 32: VRRP configuration using ACLI.....	393
Configuring global VRRP status.....	393
Assigning an IP address to a virtual router ID.....	394
Configuring the router priority for a virtual router ID.....	395
Configuring the status of the virtual router.....	395
Configuring a backup master.....	396
Configuring the critical IP address.....	396
Configuring the VRRP critical IP status.....	397
Configuring the VRRP holddown timer.....	398
Configuring VRRP holddown action.....	398
Configuring the VRRP advertisement interval.....	399
Configuring the fast advertisement interval.....	399
Configuring fast advertisement status.....	400

Configuring ICMP echo replies.....	401
Enabling VRRP traps.....	401
Displaying VRRP configuration and statistics.....	402
Chapter 33: VRRP configuration examples using ACLI.....	405
Configuring normal VRRP operation.....	405
Configuring VRRP with SMLT.....	411
Configuring VRRP with SLT.....	416
Chapter 34: VRRP configuration using Enterprise Device Manager.....	419
Configuring global VRRP status and properties.....	419
Assigning an IP address to a virtual router ID.....	421
Configuring VRRP interface properties.....	422
Graphing VRRP interface information.....	424
Viewing general VRRP statistics.....	425
Chapter 35: IGMP snooping configuration using ACLI.....	427
Configuring IGMP snooping on a VLAN.....	430
Configuring IGMP send query on a VLAN.....	431
Configuring IGMP proxy on a VLAN.....	431
Configuring the IGMP version on a VLAN.....	432
Configuring static mrouter ports on a VLAN.....	433
Displaying IGMP snoop, proxy, and mrouter configuration.....	434
Configuring IGMP parameters on a VLAN.....	435
Configuring the router alert option on a VLAN.....	437
Displaying IGMP interface information.....	437
Displaying IGMP Multicast filtering mode.....	439
Configuring IGMP multicast filtering mode.....	439
Displaying IGMP group membership information.....	440
Configuring unknown multicast packet filtering.....	441
Displaying the status of unknown multicast packet filtering.....	442
Specifying a multicast MAC address to be allowed to flood a VLAN.....	443
Displaying the multicast MAC addresses for which flooding is allowed.....	444
Displaying IGMP cache information.....	444
Flushing the router table.....	445
Configuring IGMP selective channel block.....	446
Chapter 36: IGMP snooping configuration using Enterprise Device Manager.....	449
Configuring IGMP snoop, proxy, and IGMP parameters on a VLAN using EDM.....	449
Configuring IGMP snoop, proxy, and static mrouter ports on a VLAN using EDM.....	451
Flushing the IGMP router tables and configuring IGMP router alert using EDM.....	452
Configuring unknown multicast filtering using EDM.....	455
Specifying a multicast MAC address to be allowed to flood all VLANs using EDM.....	455
Specifying an IP address to be allowed to flood a VLAN using EDM.....	456
Configuring SSM for IGMP.....	457
SSM map configuration.....	458
Displaying the SSM mapping table.....	458
Creating an SSM map for IGMP.....	459
Modifying an SSM map.....	459
Displaying IGMP cache information using EDM.....	460
Displaying IGMP group information using EDM.....	461

Displaying extended interface IGMP group information.....	462
Displaying multicast route information using EDM.....	462
Displaying multicast next hop information using EDM.....	463
Displaying multicast interface information using EDM.....	465
Creating an IGMP profile.....	465
Configuring the IGMP Profile Range.....	466
Chapter 37: PIM configuration using ACLI.....	469
Enabling and disabling PIM-SM globally.....	472
Enabling and disabling PIM-SSM globally.....	473
Configuring global PIM-SM properties.....	473
Displaying global PIM-SM properties.....	475
Enabling PIM-SM on a VLAN.....	476
Configuring the PIM-SM interface type on a VLAN.....	477
Displaying PIM-SM neighbors.....	478
Configuring PIM-SM properties on a VLAN.....	478
Displaying the PIM-SM configuration for a VLAN.....	479
Specifying the router as a candidate BSR on a VLAN.....	480
Displaying the BSR configuration.....	481
Specifying a local IP interface as a candidate RP.....	482
Displaying the candidate RP configuration.....	483
Displaying the PIM-SM RP set.....	484
Displaying the active RP per group.....	485
Enabling and disabling static RP.....	486
Configuring a static RP.....	486
Displaying the static RP configuration.....	487
Specifying a virtual neighbor on an interface.....	488
Displaying the virtual neighbor configuration.....	489
Displaying the PIM mode.....	489
Displaying multicast route information.....	490
Chapter 38: PIM-SM/SSM configuration example using ACLI.....	493
Chapter 39: PIM-SM/SSM configuration using Enterprise Device Manager.....	511
Configuring global PIM-SM/SSM status and properties.....	513
Configuring PIM-SM/SSM status and properties for a VLAN.....	514
Configuring PIM-SM/SSM VLAN properties from the IP Routing menu.....	516
Specifying the router as a candidate BSR on a VLAN interface using EDM.....	517
Setting the C-BSR priority using EDM from the VLAN menu.....	518
Setting the C-BSR priority using EDM from the IP Routing menu.....	518
Displaying the current BSR using EDM.....	519
Specifying a local IP interface as a candidate RP using EDM.....	520
Displaying the active RP using EDM.....	521
Configuring a static RP using EDM.....	521
Enabling static RP using EDM.....	522
Specifying a virtual neighbor on an interface using EDM.....	523
Displaying PIM-SM/SSM neighbor parameters using EDM.....	524
Displaying the PIM-SM RP set using EDM.....	525
Chapter 40: Basic IPv6 routing configuration using ACLI.....	527
Basic IPv6 routing configuration procedures.....	527

Configuring global IPv6 routing status.....	528
Displaying global IPv6 configuration.....	528
Configuring an IPv6 address for a VLAN.....	529
Removing the IPv6 address configuration from a VLAN.....	530
Displaying IPv6 address configuration for a VLAN.....	531
Configuring neighbor discovery prefixes.....	531
Displaying neighbor discovery prefix configuration.....	533
Configuring router advertisement.....	533
Configuring IPv6 ICMP.....	535
Configuring IPv6 static routes.....	536
Displaying IPv6 static routes.....	537
Adding static entries to the neighbor cache.....	538
Displaying the neighbor cache.....	539
Chapter 41: Basic IPv6 routing configuration using Enterprise Device Manager.....	541
Basic IPv6 routing configuration procedures.....	541
Configuring IPv6 routing and ICMP.....	541
Configuring a link-local address for a VLAN.....	542
Displaying statistics for an IPv6 interface.....	544
Configuring an IPv6 address for a VLAN.....	545
Configuring an IPv6 discovery prefix.....	546
Configuring route advertisement.....	548
Creating IPv6 static routes.....	550
Configuring the neighbor cache.....	551
Configuring the IPv4 remote access list using EDM.....	553
Configuring the IPv6 remote access list using EDM.....	554
Graphing IPv6 interface ICMP statistics using EDM.....	555
Viewing ICMP message statistics using EDM.....	556
Viewing global IPv6 TCP properties using EDM.....	556
Viewing IPv6 TCP connections using EDM.....	557
Viewing IPv6 TCP listeners using EDM.....	558
Viewing IPv6 UDP endpoints using EDM.....	558
Chapter 42: IPv6 DHCP Relay configuration using ACLI.....	561
Specifying a local DHCP relay agent and remote DHCP server.....	561
Displaying the DHCP relay configuration.....	562
Configuring DHCP relay status and parameters on a VLAN.....	563
Displaying the DHCP relay configuration for a VLAN.....	564
Displaying DHCP relay counters.....	565
Chapter 43: IPv6 Tunnel configuration using ACLI.....	567
IPv6 tunnel configuration procedures.....	567
Configuring manual IPv6-in-IPv4 tunnels using the ACLI.....	567
Displaying manual tunnel configuration.....	568
Chapter 44: IPv6 DHCP Relay configuration using Enterprise Device Manager.....	571
Configuring the DHCP relay forwarding path.....	571
Configuring DHCP relay interface parameters.....	572
Displaying DHCP Relay statistics.....	573
Chapter 45: IPv6 Tunnel configuration using Enterprise Device Manager.....	575
Configuring a tunnel for IPv6 VLANs to communicate through an IPv4 network.....	575

Viewing the local IPv6 address associated with a tunnel.....	576
Modifying tunnel hop limits.....	578
Creating IPv6 static routes.....	580

Chapter 1: New in this release

The following sections detail what's new in *Avaya Ethernet Routing Switch 5000 Series Configuration — IP Routing Protocols* (NN47200-503) for Release 6.3.

Features

See the following sections for information about feature changes:

Circuitless IP

A circuitless IP (CLIP) interface is a loopback interface that you can use to ensure connectivity to an ERS 5000 Series device, as long as a network path is available to reach the device.

For more information about CLIP, see the following sections:

- [Circuitless IP](#) on page 72
- [Accessing Loopback Interface Configuration mode](#) on page 137
- [Configuring a CLIP interface](#) on page 138
- [Deleting CLIP configuration parameters](#) on page 139
- [Restoring CLIP to default](#) on page 139
- [Displaying CLIP information](#) on page 140
- [Configuring a CLIP interface](#) on page 145
- [Deleting a CLIP interface](#) on page 146
- [Configuring a CLIP interface for OSPF](#) on page 146

Content-Based Forward to Next Hop Enhancements

Enhancements have been made to the Content-Based Forward to Next Hop feature to increase the usability and functionality of this feature:

- per VLAN enable — Release 6.3 enables filter instances to be applied to ports on a per VLAN basis. Policy instances will only be installed on ports that are members of one or more of the VLANs that are attached to the forward to next hop policy.
- Match on TCP/UDP Port — the match criteria is expanded in Release 6.3 to include Layer 4 port numbers/ranges. You can specify a UDP or TCP (or both) port number or port number range in the next hop policy.
- Support Redundant Next Hop for Same Source IP — previous releases limited the next hop to a single destination. You can now configure a secondary next hop should the primary become unavailable or unreachable.

IGMP functionality over split multilink trunking (SMLT) network topologies

To support internet group management protocol over split multilink trunking, the following basic operations are performed:

- The groups on one SMLT aggregation switch are populated in the other SMLT aggregation switch over the IST to have a redundant path (Wiring Closet Switch). IGMP IST report messages are forwarded only to query ports and not IST links. If the report is received over SMLT, the switch then configures the SMLT links as the member ports for that group.

*** Note:**

If the IST group report message does not have SMLT ID information, then IST links are configured as the member ports for the group.

- IGMP queries received on an aggregation switch are flooded to all links (except the IST links) in the VLAN. IGMP queries reaching one aggregation switch are updated in the other aggregation switch using IST query messages. The aggregation switch uses the SMLT ID in the query message to configure the SMLT links as the querier ports. If an SMLT link goes down, the querier port moves to the IST link. This is to make sure that IGMP reports are always forwarded to the IGMP querier/PIM router.

*** Note:**

If the IGMP query is received on a link or MLT (not SMLT), the SMLT ID is set to zero in the message. The aggregation switch which receives the IST query message uses the SMLT ID in the query message to configure the SMLT links as the querier ports. If the “IST query message” does not have a SMLT ID (SMLT ID = 0), then the IST link is configured as the querier port.

- IGMP leave messages received on an aggregation switch are treated in the same way as IGMP reports. Leave messages are sent to the other aggregation switch over IST links in IST leave messages. The aggregation switch that receives the IST leave message removes the group membership for that member. If a switch has no more receivers for a

group (receivers on SMLT links are exempted), switch sends an IST IGMP prune option. The switch that receives the IST IGMP leave with prune, removes the group membership of the IST link (if exists).

- IST messages are defined for different IGMP packet types. The different message types supported are:
 - IST IGMP Query message - contains the source IP address (querier IP address), VLAN, Group Address (in the case of group specific query), maximum response time and the SMLT ID (in the case IGMP queries are received on an SMLT).
 - IST IGMP Group Query message
 - IST IGMP Group Report message - contains a source IP address, group address, VLAN, and SMLT.
 - IST IGMP Group Leave message
- There are 2 types of SMLT messages that are processed by IGMP:
 - IGMP_PEER_SMLT_DOWN: The SMLT peer that receives this message sends IST IGMP group reports to the other peer for all members of that specific SMLT. The SMLT id is set to 0, so all reports are learned on the other peer's IST port (also applicable to the querier if it is located on a port member).
 - IGMP_PEER_SMLT_UP: Same as IGMP_PEER_SMLT_DOWN, however the SMLT id is set to the proper value.
- If an SMLT peer receives an SMLT IGMP message and the SMLT id specified in that message is down, it programs the IST port instead. The IST has to be up and that particular SMLT has to be configured.

IGMP Multicast flood control

IGMP Multicast flood control limits IP multicast traffic without inhibiting other control protocols. By minimizing IP multicast flooding in the network, it eliminates the necessity of queries sent by the switch when IGMP snooping is enabled. For more information on this feature see:

[IGMP Multicast flood control](#) on page 119

PIM-SSM

The Ethernet Routing Switch 5000 Series supports the source filtering capability for IGMPv3 using PIM-SSM (IGMPv3 routing mode). IGMPv3 routing mode works only with PIM-SSM and is not backward compatible with IGMPv1 or IGMPv2. Source Specific Multicast (SSM) optimizes PIM-SM by simplifying the many-to-many model. SSM uses a one-to-many model that only uses a subset of the PIM-SM features. This model is more efficient and puts less of a load on multicast routing devices.

For more information, see [PIM-SSM overview](#) on page 114.

Routing IP directed broadcasts per VLAN

Routing IP directed broadcasts per VLAN allows for the processing of broadcast packets to be identified and forwarded to destination VLAN hosts. An IP directed broadcast packet is an IP packet whose destination address is a valid broadcast address for some IP subnet. User commands affect only the final transmission of the directed broadcast on its ultimate destination subnet.

For more information see: [Routing IP directed broadcasts per VLAN](#) on page 61

VRF Lite

This release supports Virtual Routing and Forwarding (VRF) Lite for the ERS 5600 series only. With VRF Lite, you can reduce operating costs by employing virtual router instances to maintain networking capabilities and traffic isolation for clients operating over the same physical router.

For more information about VRF Lite, see the following sections:

- [VRF Lite](#) on page 70
- [VRF Lite configuration using ACLI](#) on page 149
- [VRF Lite configuration using EDM](#) on page 157

Chapter 2: Introduction

This document provides procedures and conceptual information to configure IP routing features on the Avaya Ethernet Routing Switch 5000 Series , including RIP, OSPF, VRRP, static routes, Proxy ARP, DHCP Relay, and UDP forwarding. It also provides procedures and conceptual information to manage multicast traffic using PIM-SM and IGMP snooping.

ACL I command modes

ACL I provides the following command modes:

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

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

If no password is set, you can enter ACL I 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.

Command mode and sample prompt	Entrance commands	Exit commands
User EXEC 5650TD>	No entrance command, default mode	exit or logout
Privileged EXEC 5650TD#	enable	exit or logout
Global Configuration 5650TD(config)#	From Privileged EXEC mode, enter: configure	To return to Privileged EXEC mode, enter: end

Command mode and sample prompt	Entrance commands	Exit commands
		or exit To exit ACLI completely, enter: logout
Interface Configuration 5650TD(config-if)#	From Global Configuration mode, to configure a port, enter: interface fastethernet <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 5650TD(config-router)#	From Global Configuration mode, to configure OSPF, enter: router ospf To configure RIP, enter: router rip To configure VRRP, enter: router vrrp	To return to Global Configuration mode, enter: exit To return to Privileged EXEC mode, enter: end To exit ACLI completely, enter: logout

For more information, see *Avaya Ethernet Routing Switch 5000 Series Fundamentals* (NN47200-104).

Chapter 3: IP routing fundamentals

This chapter provides an introduction to IP routing and the IP routing protocols used in the Avaya Ethernet Routing Switch 5000 Series.

IP address 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 1: IP address classifications

Class	Address range	Mask	Number of networks
A	1.0.0.0 - 127.0.0.0	255.0.0.0	127
B	128.0.0.0 - 191.255.0.0	255.255.0.0	16 384
C	192.0.0.0 - 223.255.255.0	255.255.255.0	2 097 152
D	224.0.0.0 - 239.255.255.254		
E	240.0.0.0 - 240.255.255.255		

*** Note:**

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

*** Note:**

Class D addresses are primarily reserved for multicast operations, although Open Shortest Path First (OSPF) uses the addresses 224.0.0.5 and 224.0.0.6, and Routing Information Protocol (RIP) uses 224.0.0.9.

*** Note:**

Class E addresses are reserved for research purposes.

To express an IP address in dotted-decimal notation, convert each octet of the IP address to a decimal number and separate them 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.

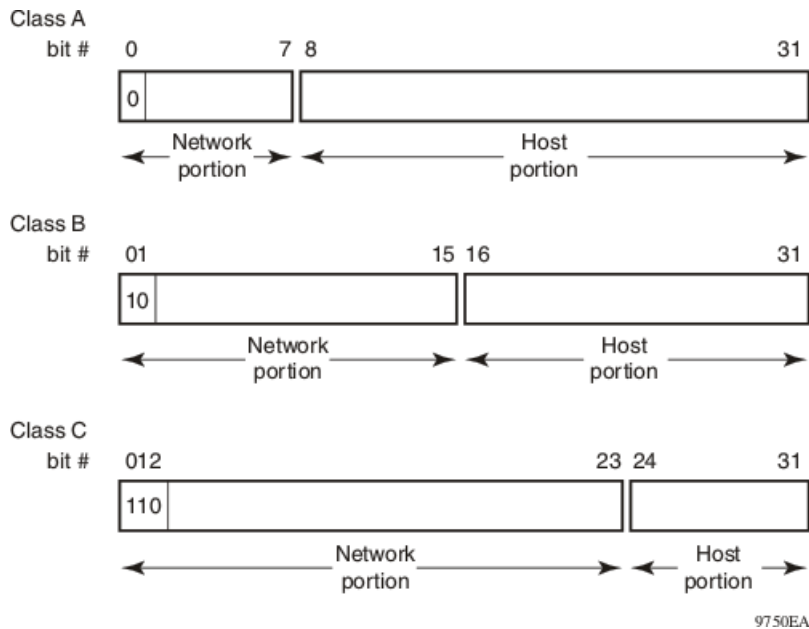


Figure 1: Network and host boundaries in IP address classes

Subnet addresses

Subnetworks (or subnets) are an extension of the IP addressing scheme. Organizations use subnets to 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).

Create a subnet address 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 applies 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 Avaya Ethernet Routing Switch 5000 Series supports.

Table 2: 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

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts per subnet
3	255.255.224.0	6	8 190
4	255.255.240.0	14	4 094
5	255.255.248.0	30	2 046
6	255.255.252.0	62	1 022
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	1 022	62
11	255.255.255.224	2 046	30
12	255.255.255.240	4 094	14
13	255.255.255.248	8 190	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) divides 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 5000 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 *Avaya Ethernet Routing Switch 5000 Series Configuration — VLANs, Spanning Tree, and Link Aggregation*, NN47200-502.

IP routing using VLANs

The Avaya Ethernet Routing Switch 5000 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 does not use a specific port association because the VLAN IP address is reachable 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.

A one-to-one correspondence does not exist between the physical port and the routable interface, because a given port can belong to multiple routable VLANs.

When the Avaya Ethernet Routing Switch 5000 Series routes IP traffic between different VLANs, the switch runs in Layer 3 mode; otherwise, the switch runs in Layer 2 mode. After you assign an IP address to a Layer 2 VLAN, the VLAN becomes a routable Layer 3 VLAN. You can assign a unique IP address to each VLAN.

You can configure the global status of IP routing to enabled or disabled on the Avaya Ethernet Routing Switch 5000 Series. By default, IP routing is disabled.

You can configure all IP routing parameters on the Avaya Ethernet Routing Switch 5000 Series before you actually enable routing on the switch.

In this release, the Avaya Ethernet Routing Switch 5000 Series supports local routes, static routes, and dynamic 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. With dynamic routing, routes are identified using a routing protocol such as RIP or OSPF.

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 that you assign to a VLAN interface, the Ethernet Routing Switch adds a directly connected or local route to its routing table based on the IP address and 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

VLAN 100 uses IP address 10.100.1.1/24 and VLAN 200 uses IP address 10.200.1.1/24. As you enable IP routing, 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, the Ethernet Routing Switch can reach both hosts A (10.200.1.10) and B (10.100.1.10). However, to achieve Layer 3 connectivity between A and B, you must perform additional configuration. 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.

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

Local and 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. In this release, the Ethernet Routing Switch supports local and non-local static routes. 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 multiple paths exist to a network and you can reduce the number of static routes 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 fails.

Default routes

Default routes specify a route to all networks for which there are no explicit routes in the Forwarding Information Base (FIB) or the routing table. The 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. The switch forwards traffic to the next-hop IP address that the default route specifies.

Management VLAN

If you enable IP routing 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, including Telnet, Web, Simple Network Management Protocol (SNMP), BootP, and Trivial File Transfer Protocol (TFTP). If you do not enable routing, you can reach the management VLAN 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. You cannot remove the management VLAN.

After you enable routing on the Avaya Ethernet Routing Switch 5000 Series switches, the management VLAN is similar to other routable VLANs. You can reach the IP address 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. The management route does not carry inter-VLAN routed traffic from the other Layer 3 VLANs to the destination network, which 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 management route next-hop IP address must reside in the same network as the IP address of the switch or stack, and it must be directly reachable.

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

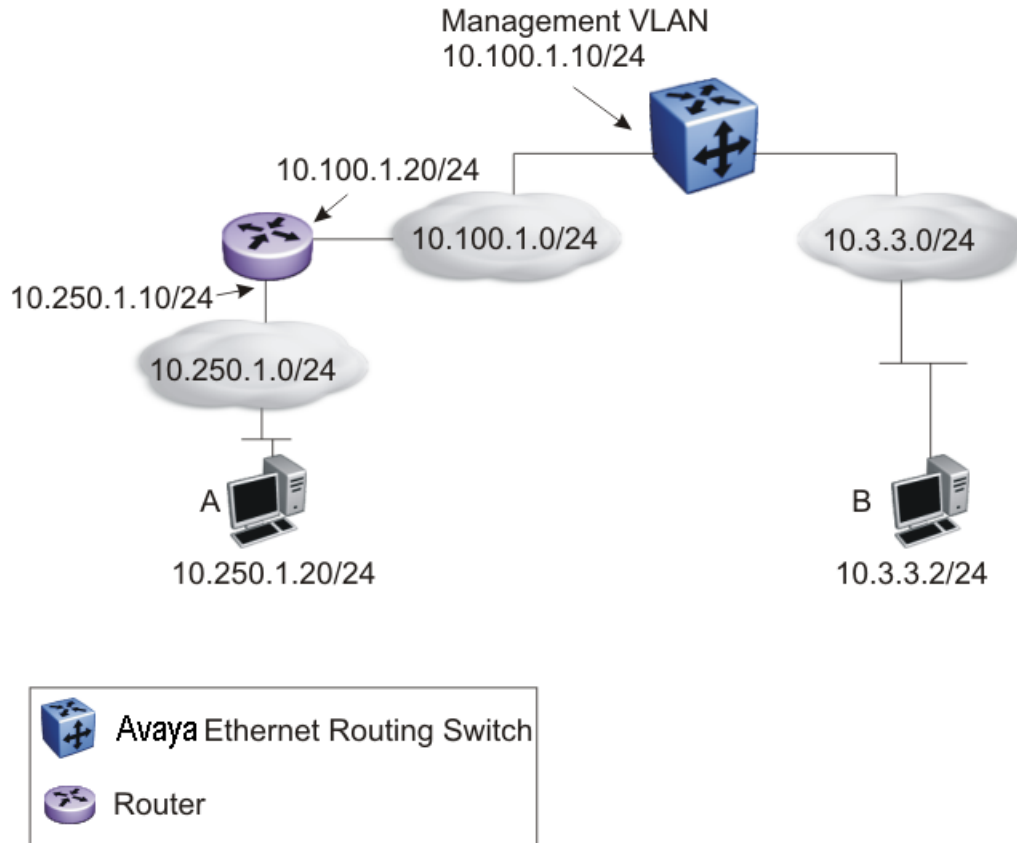


Figure 3: Management route

As network 10.250.1.0/24 does not directly connect to the Ethernet Routing Switch, to achieve connectivity from host 10.250.1.20 to the management VLAN, you must configure the Ethernet Routing Switch 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 perform management operations, 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 you configure a Layer 3 VLAN for network 10.3.3.0/24, you provide 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.

Multinetting

The Avaya Ethernet Routing Switch 5000 Series supports the definition and configuration of up to eight secondary interfaces on each VLAN (multinetting). With IP multinetting, you can associate multiple IP subnets with one VLAN. That is, connected hosts can belong to different IP subnets on the same VLAN.

You can configure multinetting using ACLI or Enterprise Device Manager.

The following diagram illustrates a network with IP multinetting.




Figure 4: Network with Multinetting

You can configure a static route with the next hop on the secondary interface. You can also add static Address Resolution Protocol (ARP) for an IP address in the same subnet of a secondary interface.

The following list provides limitations for secondary interfaces:

- You can have a maximum of eight secondary interfaces on each VLAN.
- You can have a maximum of 256 IP interfaces (including primary and secondary).
- You enable or disable all secondary interfaces on a VLAN simultaneously. You cannot configure the administrative state of the secondary IP interfaces individually.

- Dynamic routing is not available for secondary IP interfaces.
- Routers do not support secondary interfaces.
- A primary IP interface must exist before you can add secondary IP interfaces; you must delete secondary interfaces before you can delete the primary interface.

If you configure secondary interfaces on the management VLAN, you cannot disable routing globally or on the management VLAN. NVRAM purges secondary IP interfaces on the management VLAN after the following actions occur:

- a unit leaves the stack and the switch does not have a manually configured IP address
- the switch fails to obtain the IP address through the BootP mode

Secondary interfaces do not support the following protocols or features:

- Dynamic Host Configuration Protocol (DHCP)
- Proxy ARP
- User Datagram Protocol (UDP) broadcast
- IPFIX
- Virtual Router Redundancy Protocol (VRRP)
- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Border Gateway Protocol (BGP)

Brouter port

The Avaya Ethernet Routing Switch 5000 Series supports the configuration of brouter ports. 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 still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol (STP) recalculations in routed traffic. A brouter port is actually a one-port VLAN; therefore, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

When you create a brouter port, the switch also performs the following actions:

- creates a port-based VLAN
- adds the brouter port to the new port-based VLAN
- changes the PVID of the brouter port to the VLAN ID of the new VLAN

- disables the STP participation of the brouter port
- assigns an IP address to the brouter VLAN

RIP

RIP is a standards-based, dynamic routing protocol based on the Bellman-Ford (or distance vector) algorithm. RIP is an Interior Gateway Protocol (IGP). Routers use RIP to exchange information to compute the shortest routes through an IPv4-based network. RIP uses the hop count 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).

RFC 1058 defines RIP version 1 and RFC 2453 defines 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, routers do not exchange mask information; the router that receives the update always applies the natural mask. For RIP version 2, routers always include mask information.

RIP uses UDP data packets to exchange routing information.

The following list explains the sequence of processes governed by RIP:

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. All interfaces configured to send routing information send RIP advertisements.
3. The neighbors send their routing tables, and then 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 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

one hop. This cost, or hop count, is the metric. The following figure shows the hop count between various units in a network.




Figure 5: RIP hop count

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.

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, routers do not exchange mask information; the router that receives the update always applies the natural mask. With RIP version 2, routers always include mask information.

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 the router receives a valid update for a garbage route within the holddown period, the router adds the route back into its routing table. If the router does not receive an update, the router completely deletes all garbage list entries for the nonupdating router.

Split horizon

To prevent routing loops, RIP uses the mechanism of split horizon, with or without poison reverse. Simple split horizon means that RIP does not advertise IP routes learned from a neighbor, back in updates to that neighbor. Split horizon with poison reverse means that RIP does advertise these routes back to the neighbor, but they are poisoned with a metric of 16, which represents infinite hops in the network. The receiver neighbor therefore ignores this route.

By default, RIP split horizon is enabled without poison reverse.

Triggered updates

To promote fast convergence, RIP also supports a triggered updates option. With triggered updates enabled, a router sends update messages whenever it changes the metric for a route, even if it is not yet time for a regular update message.

RIP send and receive modes

You can configure RIP to use a number of different send and receive modes depending on the specific network configuration.

The following table lists the send modes supported.

Table 3: RIP send modes

Send mode	Description	Result
rip1comp	This mode broadcasts RIP version 2 updates using RFC 1058 route consumption rules. This mode is the default send mode for the Avaya Ethernet Routing Switch 5000 Series .	<ul style="list-style-type: none"> • Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 192.1.2.255) • RIP Update is formed as a RIP version 2 update, including network mask • RIP version = 2

Send mode	Description	Result
rip1	This mode broadcasts RIP updates that are compliant with RFC 1058.	<ul style="list-style-type: none"> • Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 192.1.2.255) • RIP Update is formed as a RIP version 1 update, no network mask included • RIP version = 1
rip2	This mode broadcasts multicast RIP version 2 updates.	<ul style="list-style-type: none"> • Destination MAC is a multicast, 01-00-5e-00-00-09 • Destination IP is the RIP version 2 multicast address, 224.0.0.9 • RIP Update is formed as a RIP version 2 update including network mask • RIP version = 2
nosend	No RIP updates are sent on the interface.	None

The following table lists the receive modes supported.

Table 4: RIP receive modes

Receive mode	Result
rip1OrRip2	The switch accepts RIP version 1 or RIP version 2 updates.
rip1	The switch accepts RIP version 1 and RIP version 1 compatible updates.
rip2	The switch accepts RIP version 2 updates.

Supported RIP capabilities on the 5000 switch

RIP supports the following standard behaviors:

- periodic RIP updates about effective best routes
- garbage collection
- split horizon with or without poison reverse
- triggered update for changed RIP routes
- unicast to the specific query requestor
- broadcast and 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 for each interface
- cost for each device and for each interface
- equal cost multipath (ECMP)

The Avaya Ethernet Routing Switch 5000 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

You can configure many RIP features. The actual behavior of the protocol depends on the feature configuration.

RIP limitations

RIP has the following limitations:

- The protocol is limited to networks with a longest path of 15 hops.
- The protocol depends on counting to infinity to resolve certain unusual situations.
- The protocol uses fixed metrics (the hop number) to compare alternative routes, as opposed to real-time parameters such as measured delay, reliability, or load.
- RIP does not support address-less links.

OSPF

OSPF is a classless IGP that distributes routing information between routers that belong 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 VLSM and tagging of externally-derived routing information.

Important:

The Avaya Ethernet Routing Switch 5000 Series implementation of OSPF only supports broadcast and passive interfaces. Point-to-point and NBMA interfaces are not supported.

Overview

In an OSPF network, each router maintains a link-state database that describes the topology of the 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:

- Fast convergence

After the network topology changes, OSPF recalculates routes quickly.

- Minimal routing protocol traffic

Unlike distance vector routing protocols, such as RIP, OSPF generates a minimum of routing protocol traffic.

- Load sharing

OSPF supports ECMP routing. If several equal-cost routes to a destination exist, traffic is distributed equally among them.

- Scalable

Because OSPF does not use hop count in its calculation, the routing domain is scalable.

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

The Ethernet Routing Switch 5000 Series implementation of OSPF does not support TOS-based routing.

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, you can divide an OSPF network 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 you divide a network 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 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 it does not use routing information obtained from outside the area.

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 that attaches to two or more areas inside an OSPF network is 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 you divide the AS into multiple areas, each nonbackbone area must attach 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.

ABRs 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. 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 divided into three areas: a backbone area, a stub area, and a not-so-stubby area (NSSA). This document describes stub areas and NSSAs in [Area types](#) on page 42.

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.

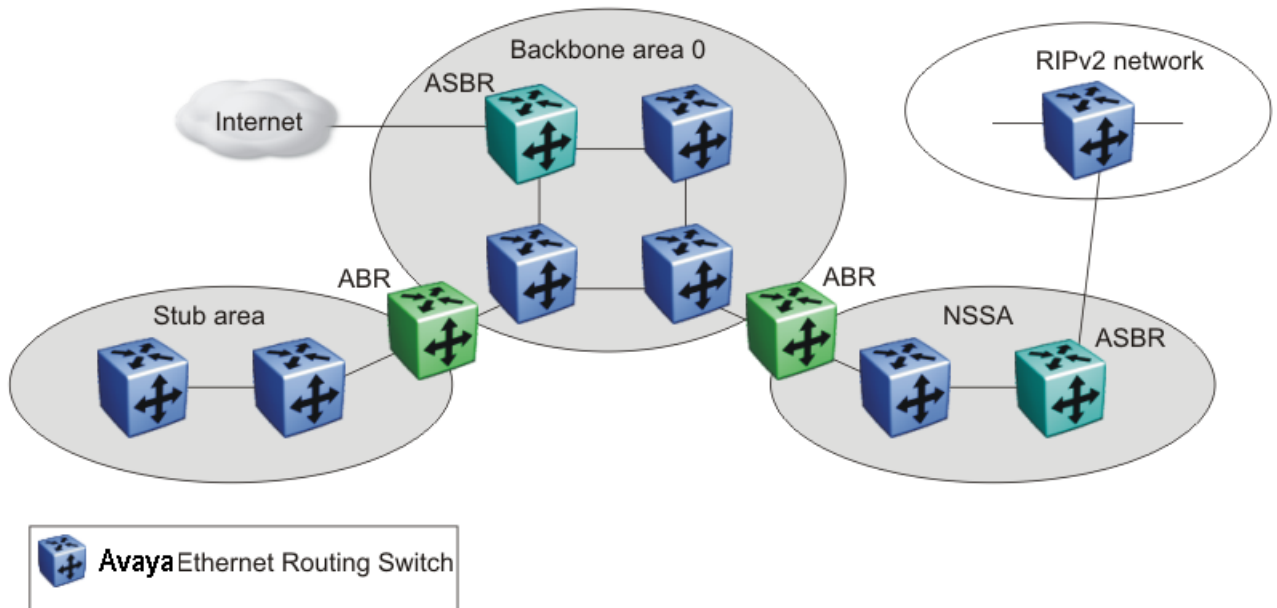


Figure 6: OSPF network

ASBR and external router advertisements

A router functions as an ASBR if one or more of its interfaces connect to a non-OSPF network, for example, RIP, or static routes. ASBRs advertise non-OSPF routes into OSPF domains.

An ASBR advertises external routes into the OSPF domain using autonomous system external (ASE) LSAs (LSA type 5). ASE LSAs flood across area borders. To conserve resources or control traffic flow, you can limit the number of ASBRs in your network.

OSPF considers the following routes to be ASE routes:

- A route to a destination outside the AS
- A static route
- A default route
- A route derived by RIP
- A directly connected network that does not run OSPF

External route metrics

When an ASBR imports external routes, it imports OSPF route information using external type 1 or type 2 metrics. With external type 1 metrics, OSPF calculates the total cost by adding the

external metric value and the internal cost to the ASBR. For external type 2 metrics, OSPF uses only the internal OSPF cost to the ASBR in the routing decision. You can specify the metric type to use by configuring a route policy. For more information, see [Route policies](#) on page 52.

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. After the routers synchronize their databases, the routers are fully adjacent.

To limit the amount of routing protocol traffic, OSPF routers use the Hello protocol to elect a DR and a 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

The following list defines the sequence of OSPF processes on broadcast multiaccess networks:

1. After 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 DR and 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 that contain 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 converges, each OSPF router continues to periodically flood hellos to maintain neighbor relationships. 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). After the DR receives 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 router advertisements

An OSPF router advertisement expresses a destination 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

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 5: Router types in an OSPF network

Router type	Description
AS boundary router (ASBR)	A router that attaches at the edge of an OSPF network is an ASBR. Any router that distributes static routes or RIP routes into OSPF is

Router type	Description
	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 that attaches to two or more areas inside an OSPF network is 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 exist.
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 synchronize and advertise 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, the BDR can assume the DR role quickly.

LSA types

After the network converges, OSPF does not require each router to send 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. OSPF routers generate LSAs either periodically, to ensure connectivity, or 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 6: OSPF LSA types

LSA type	LSA name	Description	Area of distribution
1	Router LSA	Every router generates Type 1 LSAs to describe their set of active interfaces and neighboring routers. Routers flood Type 1 LSAs 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 DR generates network LSAs that list all routers on that LAN. Routers flood Type 2 LSAs only within the area. A backbone DR can flood network links	Only within the same area

LSA type	LSA name	Description	Area of distribution
		advertisements within the backbone area.	
3	Network-Summary LSA	The ABR generates Type 3 LSAs to describe the networks that are reachable outside the area. An ABR that attaches to two areas generates a different network summary LSA for each area. ABRs also flood Type 3 LSAs that contain information about destinations within an area to the backbone area.	Passed between areas
4	ASBR-summary LSA	The ABR generates Type 4 LSAs to advertise the cost of the path to the closest ASBR from the router that generates the advertisement.	Passed between areas
5	Autonomous System External [ASE] LSA	The ASBR generates Type 5 LSAs to describe the cost of the path to a destination outside the AS from the ASBR that generates the advertisement. Routers pass Type 5 LSAs 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	OSPF NSSAs use Type 7 LSAs 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

The following figure shows a stub area at the edge of the OSPF routing domain with only one ABR.

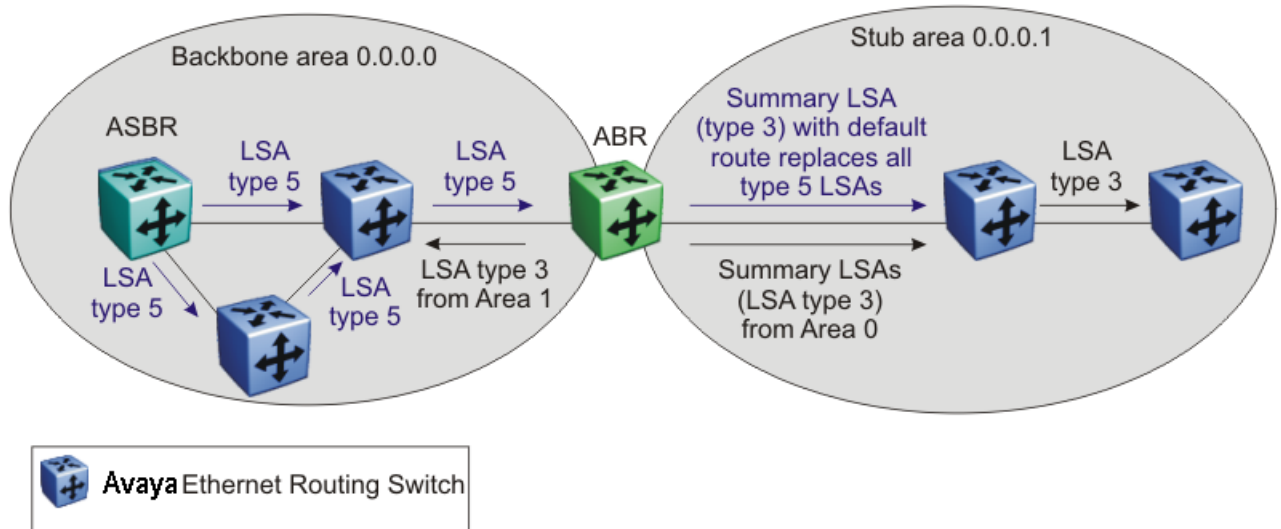


Figure 7: Stub area

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. Because 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, 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.

Totally stubby area

To further reduce the size of the stub area LSDB, you can configure a totally stubby area, which prevents redistribution of summary routes (Summary LSAs, type 3) from other areas into the stub area. As shown in the following figure, the totally stubby area ABR advertises a default route into the stub area not only for external routes, but for all destinations outside of the stub area.

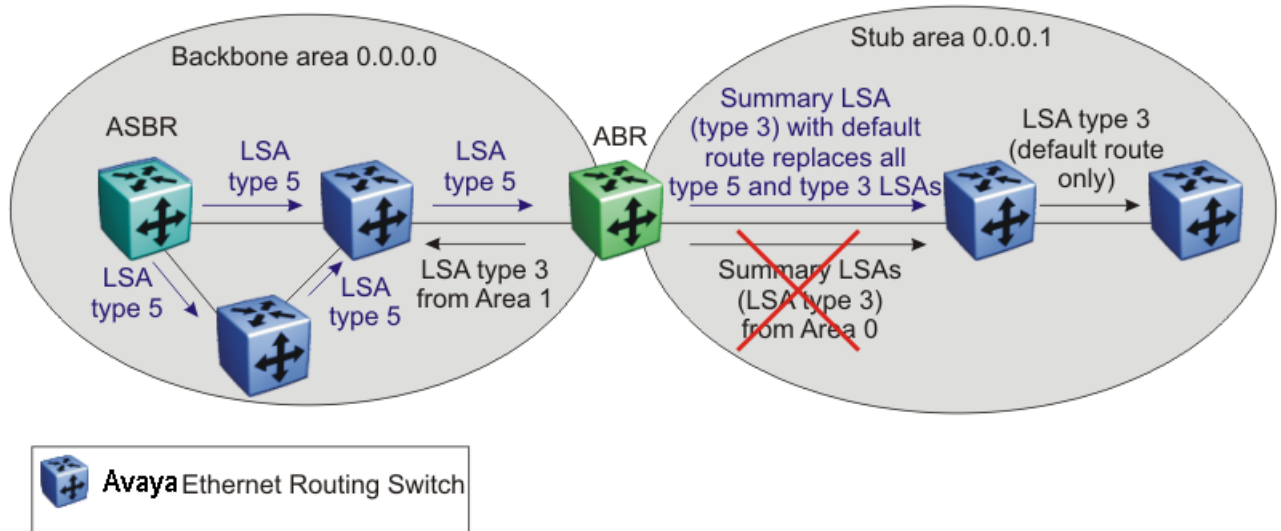


Figure 8: Totally stubby area

To configure a totally stubby area, you must disable import summaries on the stub area ABR.

You can only disable import summaries in the stub area.

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. The NSSA can 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.

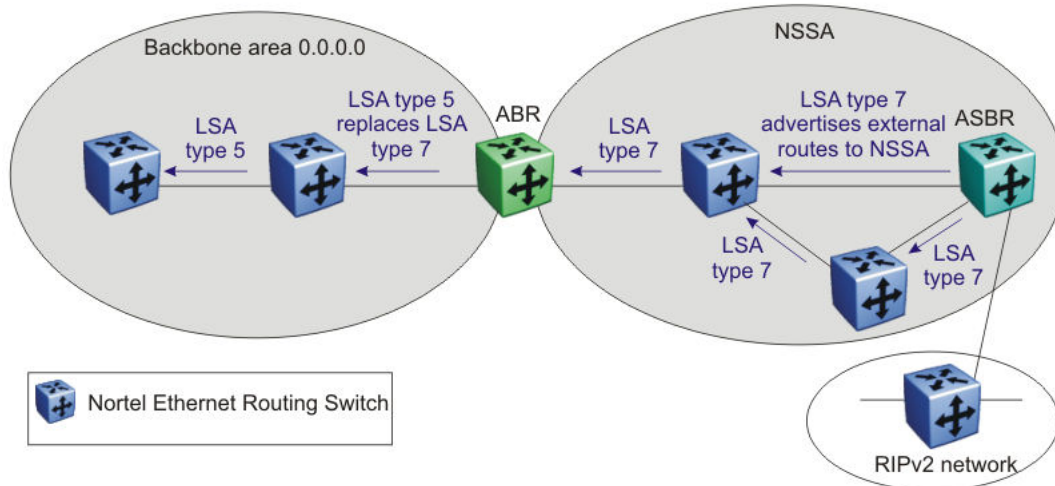


Figure 9: OSPF NSSA

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 flooding the external routes to other areas. To support this additional control over external router advertisement, the type 7 LSAs provide an Options field that contains an N/P-bit that notifies the ABR which external routes can be advertised to other areas. When the NSSA N/P-bit is true (the default configuration), 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 5000 Series .

Normal area

A normal area is an area that is neither a backbone nor a stub area that sends and receives LSA types 1 through 5. The following figure shows a normal area that supports ABRs and ASBRs.

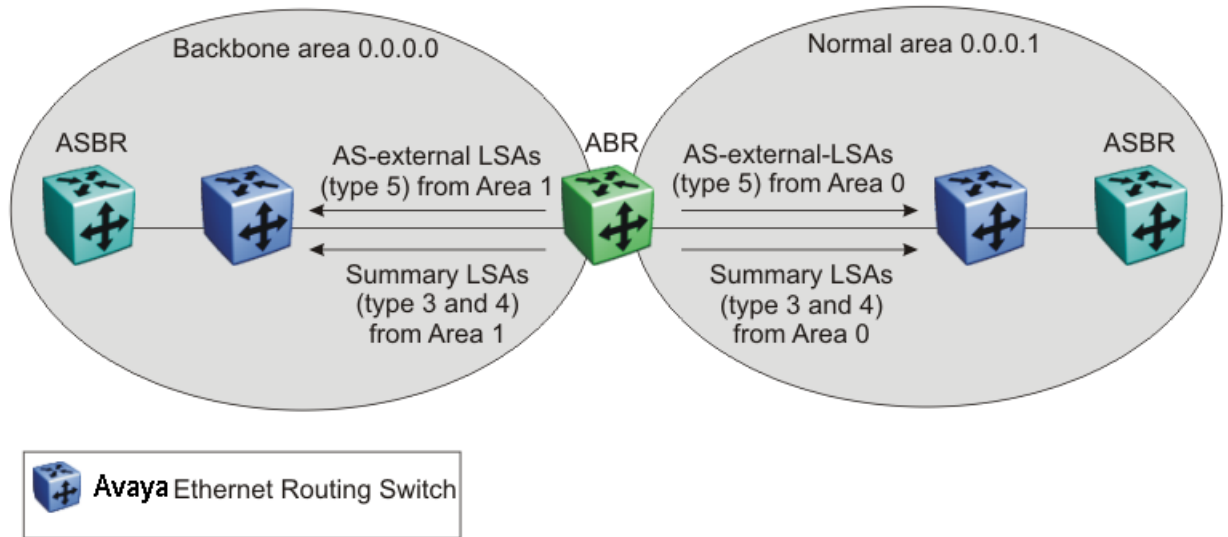


Figure 10: OSPF normal area

The Avaya Ethernet Routing Switch 5000 Series automatically becomes an ABR when it connects to more than one area.

OSPF virtual link

You can partition the OSPF network into multiple areas. However, every non-backbone area must connect 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 11: 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 that provides 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 if you delete the transit area, disable auto virtual link, 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 exist 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 you cannot modify. 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.

Identify host routes 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.

After you add a host to, or delete from, a host route, the router updates the router LSAs and floods them to neighbors in each area where that router has an interface.

The following is an example of parameters for a host route advertised in the LSA.

Host route in LSA

- Type: 3 (stub network)
- LinkID: IP address of host directly connected to router
- Link Data: 0xFFFFFFFF
- Metric: configured cost of host

OSPF interfaces

Configure an OSPF interface, or link, on an IP interface. In the Ethernet Routing Switch 5000, an IP interface can be either a brouter port or a VLAN. The state information associated with the interface is obtained 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 reenabling 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. The following sections describe the OSPF network interface types that the Ethernet Routing Switch 5000 supports.

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

Configure the interface as passive to advertise it as an OSPF internal route. If the interface is not a passive interface, to advertise a network into OSPF and not form OSPF adjacencies, you must configure the interface as nonOSPF, and redistribute the local network as an autonomous system external (ASE) LSA.

OSPF treats the network behind a passive interface 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 that sends the packet. An OSPF packet is one of the following types:

- The router transmits hello packets between neighbors and never forwards them. The Hello protocol requires routers to send hello packets to neighbors at predefined hello

intervals. A neighbor router that does not receive a hello packet declares the other router dead.

- The router exchanges database description (DD) packets after neighboring routers establish a link, which synchronizes their LSDBs.
- Link-state request packets describe one or more link-state advertisements that a router requests from its neighbor. A router sends 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 the router sends them following a change in network conditions.
- The router sends link-state acknowledgement packets to acknowledge receipt of linkstate updates. Link-state acknowledgement packets contain the headers of the received LSAs.

OSPF metrics

For OSPF, the best path to a destination is the path that offers the least-cost metric (least-cost delay). You can configure OSPF cost metrics to 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.

The following table shows the default metric speeds for different port types.

Table 7: OSPF default metrics

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

Automatic router ID change in a stack

If a unit leaves the stack and becomes standalone (when the stack disjoins), the router ID automatically changes to its default value. This change prevents router ID duplication in the OSPF routing domain.

For this feature to operate, you must turn off IP blocking (set to none), and globally enable OSPF.

The new router ID value is temporary, that is, it is not saved in NVRAM. Therefore, upon reset, the old router ID is restored.

OSPF security mechanisms

The Ethernet Routing Switch 5000 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 an incorrectly-configured router from joining an OSPF domain. The switch supports two security mechanisms: 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:

Do not 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

Use 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 match between sending and receiving routers. Both the sending and receiving routers calculate the message digest based on the MD5 key and any padding, and then compare the results. If the message digest computed at the sender and receiver does not match, the packet is rejected.

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

Equal Cost MultiPath

Routers use the Equal Cost MultiPath (ECMP) feature to determine equal cost paths to the same destination prefix. Routers can use the multiple paths for load sharing of traffic and the multiple paths provide faster convergence to other active paths in case of network failure. By

maximizing load sharing among equal-cost paths, routers use links more efficiently when sending IP traffic. The ECMP feature supports the following protocols:

- OSPF
- RIP
- Border Gateway Protocol (BGP)
- Static routes

Avaya Ethernet Routing Switch 5510 does not support ECMP. ECMP works in a mixed stack but cannot run on Avaya Ethernet Routing Switch 5510 units in the stack.

Route policies

Using standard routing schemes, a router forwards packets on routes that it learns 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 that the network administrator creates. The administrator then applies these rule sets, or policies, to the learned or static routes.

On the Avaya Ethernet Routing Switch 5000 Series, you can configure route policies for RIP and OSPF. When used in conjunction with these protocols, route policies can perform the following tasks that are not possible using traditional routing methods:

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

The Ethernet Routing Switch 5000 Series supports the following types of policies:

- Accept (In) policies

Accept policies apply to incoming routing updates before they are applied to the routing table. In the case of RIP and BGP, you can apply accept policies to all incoming packets. You can create only one policy for each RIP interface. In the case of OSPF, accept policies only apply to Type 5 External routes based on the advertising router ID. Only one OSPF accept policy can exist for each switch, and the policy applies before updates are added to the routing table from the link state database.

- Announce (Out) policies

Announce policies apply to outgoing routing updates before the switch transmits routing update packets. In the case of RIP, you can apply announce policies to all outgoing packets. You can create only one policy for each RIP interface. OSPF does not support

announce policies because OSPF requires routing information to be consistent throughout the OSPF domain.

- Redistribution policies

Use redistribution policies 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. Only one OSPF redistribution route can exist for each switch, and you must enable redistribution. 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.
- Identified by a prefix list name and unique identifier.
- Prefix lists support the comparison of ranges of incoming masks.

- Route maps

- Contain a set of match and set parameters.
- Match and set parameters can contain several prefix lists.
- A sequence number identifies a set of match and set parameters.
- Accept and deny actions are associated with each sequenced parameter set.
- Sequence numbers act as a preference configuration. Sets with a lower sequence number are preferred over those with a higher sequence number.

To configure routing policies, create the appropriate prefix lists, and then assign those prefix lists to route maps. After you create all the route maps, assign them to the appropriate type of policy.

Route policies in a stack

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

- All stack units store the policy database.
- Only the base unit supports policy configuration. 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 synchronizes with the non-base unit. The database in the non-base units are deleted during the exchange.
- RIP and OSPF use only the policies in the base unit for policy application.

DHCP relay

Dynamic Host Configuration Protocol (DHCP) assigns network IP addresses on a dynamic basis to clients that request an address. DHCP is an extension of the Bootstrap protocol (BootP). BootP and DHCP clients (workstations) generally use 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 5000 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. The switch also relays server replies back to the clients.

To relay DHCP messages, you must create two Layer 3 VLANs: one that connects to the client, and another that provides a path to the DHCP server. You can enable DHCP Relay on an individual VLAN basis.

DHCP-relay support exists on all VRF instances. You can configure the forwarding path while in VRF Router Configuration mode for non-default VRFs. See [Entering Router Configuration mode](#) on page 137 and [DHCP relay configuration using ACLI](#) on page 349.

! Important:

The DHCP Relay feature shares resources with Quality of Service (QoS). If you enable the DHCP Relay feature, you cannot install a QoS policy with a precedence of 11.

For more information on QoS policies, see *Avaya Ethernet Routing Switch 5000 Series Configuration - Quality of Service* (NN47200-504).

The following figure shows a DHCP Relay example. An end station connects to subnet 1, which corresponds to VLAN 1. The Avaya Ethernet Routing Switch 5000 Series connects two subnets by using the virtual routing function. If 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), and the DHCP Relay is enabled, the Ethernet Routing Switch forwards DHCP requests to the host address of the DHCP server on VLAN 2.




Figure 12: 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 5000 Series to forward DHCP packets from the end station to the server, use 10.10.2.1 as the server address.

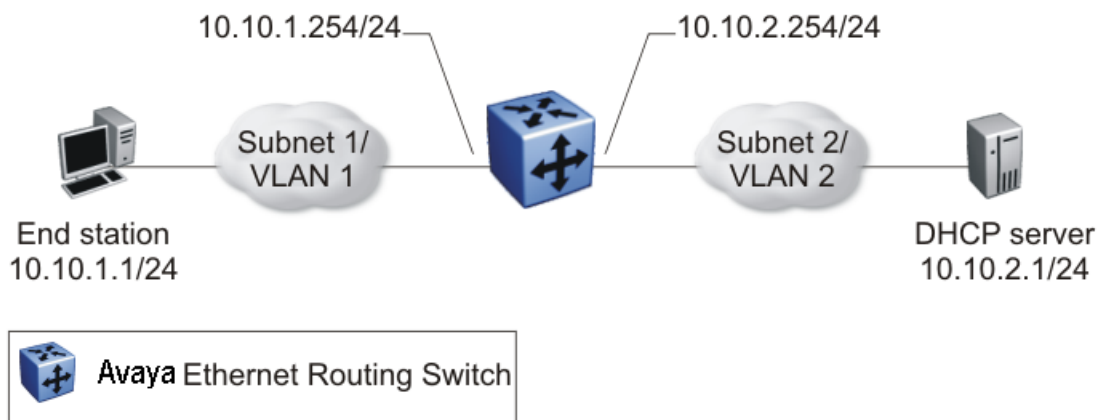


Figure 13: Forwarding DHCP packets

The switch forwards all BootP and DHCP broadcast packets that appear on the VLAN 1 router interface (10.10.1.254) to the DHCP server. In this case, the switch forwards the DHCP packets 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 5000 Series can forward DHCP requests to multiple servers. You can configure up to 512 servers to receive copies of the forwarded DHCP messages.

To configure DHCP client requests forwarding 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 on different VLANs. To configure the Avaya Ethernet Routing Switch 5000 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 then associate this relay agent with each of the DHCP server addresses, 10.10.2.1 and 10.10.3.1.

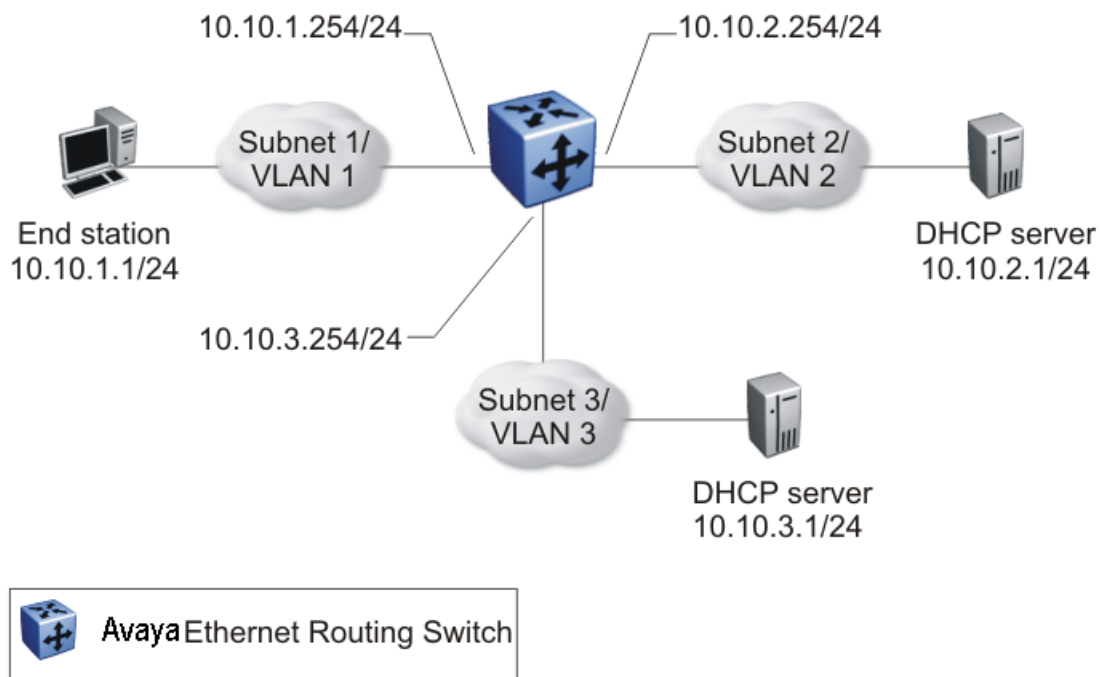


Figure 14: Multiple BootP/DHCP servers

Differences between DHCP and BootP

With DHCP relay, the Avaya Ethernet Routing Switch 5000 Series supports the relay of DHCP and the Bootstrap protocol (BootP). RFC 2131 specifies the following differences between DHCP and BootP:

- BootP enables the retrieval of an 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 needed 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 called options (RFC 2131).

IP forwarding next-hop

After a router receives a packet it normally decides where to forward it based on the destination address in the packet, which it then uses to look up an entry in a routing table. However, in some cases, there can be a need to forward a packet based on other criteria. For example, a network administrator can choose to forward a packet based on the source address, not the destination address.

The following figure shows an example in which this feature proves useful.

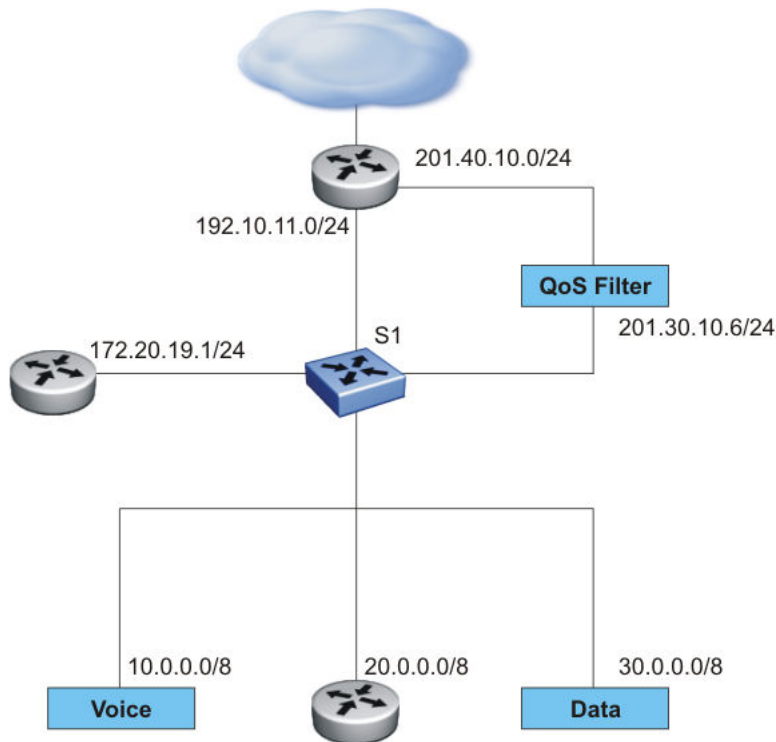


Figure 15: IP forwarding next-hop example

In the preceding figure, to reach external networks with normal routing, all traffic from subnets 10.0.0.0/8, 20.0.0.0/8, and 30.0.0.0/8 flows through 192.10.11.0/24, the best route in the S1 routing table to the outside world.

In this example, data traffic going to external networks must be directed to a filter to apply QoS parameters on the traffic. To that end, you can apply IP forwarding next-hop on S1, specifying 30.0.0.0/8 as the source address and 201.30.10.6 as the next-hop. This configuration allows the router to route the data traffic to the desired QoS filter.

Similarly, to route traffic from 10.0.0.0/8 to a different location, you can specify 10.0.0.0/8 as the source address and 172.20.19.1 as the next-hop. With this configuration, the router routes traffic from the specified subnet to the desired next hop.

Limitations

The following are the limitations for IP Forwarding Next-hop:

- The IP Forwarding Next-hop feature is supported on the Ethernet Routing Switch 5600 only.
- In a hybrid stack, you can use IP Forwarding Next-hop functionality only on the Ethernet Routing Switch 5600 unit.
- Multipath is not supported.

Enhancements

Following are the enhancements to IP Forward Next-Hop beginning with Release 6.3:

- Per-VLAN enable and disable support

You can temporarily alter the policies that apply at specific ports by administratively enabling or disabling IP Forward Next-Hop policy data that applies to a VLAN. You can administratively enable or disable all the policies associated with a VLAN with ACLI or EDM support in the Interface configuration mode.

- Port range support

You can identify Layer 4 destination ports to use for matching purposes. You can specify a single port or a port range, in addition to the mandatory source IP data, in the ip-fwd-nh policy specification. You can further constrain matching based on port type (TCP/UDP) if necessary. You cannot associate multiple port ranges with a single ip-fwd-nh policy.

For example, you can define an ip-fwd-nh policy configuration including port data as follows:

```
ip fwd-nh policy polWithPort match source-ip 10.10.10.0/24 port-min 67 port-max 80 set next-hop 10.11.11.23
```

This policy matches IPv4 traffic with the source IP subnet 10.10.10.0/24, with the IP protocol equal to TCP or UDP (port type defaults to 'both'), and with the Layer 4 destination port in the range of 67 to 80 (inclusive).

- Secondary next-hop IP address support

To improve feature flexibility, ip-fwd-nh policy support is augmented to allow you to associate a secondary next-hop IP address with a policy entry. If the primary next-hop IP address is not currently "resolved", the secondary next-hop IP address will be used if it is "resolved". A next-hop IP address is "resolved" or considered active if a physical (MAC) address is associated with the Layer 3 address through user configuration (that is a static address configuration) or through system operation (that is address resolution through ARP).

The primary next-hop IP address will always take precedence if and when it is resolved. If the secondary next-hop IP address is used for traffic forwarding when the primary next-hop IP address becomes active, the primary next-hop IP address replaces the secondary next-hop in the appropriate ip-fwd-nh traffic forwarding operations.

Most restrictions that apply to primary next-hop IP address also apply to the secondary next-hop IP address. These restrictions include:

- A broadcast address (all 1's) is not allowed
- The address must be directly reachable based on the current configuration

- Addresses that are associated with any system interfaces (the system IP addresses) are not allowed

- Filter usage optimization

Initially ip-fwd-nh policy instances applied to all ports in the system regardless of the VLAN association of the port, and resources were required on all ports for a successful installation. With Release 6.3, the ip-fwd-nh filter usage has been enhanced such that filter resources are used based on the VLAN membership of the port. Policy instances are installed only on ports that are members of one or more VLANs that are attached to the ip-fwd-nh policy.

Improved filter usage efficiency requires increased interaction with the VLAN module. Port VLAN assignments and VLAN activation or deactivation are real-time inputs to the ip-fwd-nh functionality. Filter resources apply on a port if all of the following conditions are true:

- The port is associated with VLAN X.
- VLAN X is attached to an ip-fwd-nh policy.
- VLAN X is active (VLAN interface is routing-enabled).
- VLAN X is administratively enabled.
- The ip-fwd-nh feature is enabled.

UDP broadcast forwarding

By default, the switch does not route UDP broadcast frames received on one VLAN to another VLAN. However, some network applications, such as the NetBIOS name service, rely on UDP broadcasts to request a service or locate a server. 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 router interface receives a UDP broadcast, the broadcast must meet the following criteria to be considered for forwarding:

- The broadcast must be a MAC-level broadcast.
- The broadcast must be an IP-limited broadcast.
- The broadcast must be for a configured UDP protocol.
- The broadcast must have a TTL value of at least 2.

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

When you enable the UDP forwarding feature, a filter is installed that compares the UDP destination port of all packets against all 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.

Directed broadcasts

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

If you enable IP directed broadcasting on a VLAN, the Ethernet Routing Switch forwards direct broadcast packets in 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.

Routing IP directed broadcasts per VLAN

Routing IP directed broadcasts for each VLAN allows for the processing of broadcast packets to be identified and forwarded to destination VLAN hosts. An IP directed broadcast packet is an IP packet whose destination address is a valid broadcast address for some IP subnet. User commands affect only the final transmission of the directed broadcast on its ultimate destination subnet.

When an IP directed broadcast packet is sent, the network forwards it the same way as a unicast packet. When the packet reaches a switch directly connected to the target subnet, the switch checks whether the IP directed broadcast feature is enabled both globally and on the interface that directly connects to the target subnet. If you enable IP directed broadcast on the interface, the switch broadcasts the packet on that subnet by rewriting the destination IP address as the configured broadcast IP address for the subnet. The switch converts the packet to a link layer broadcast packet that every host on the network processes. If you disable the IP directed broadcast feature, the switch drops the packet.

You can enable or disable this feature for each VLAN interface and globally. By default, the feature is disabled globally and for each VLAN interface.

ARP

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

Network stations that use 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, the network station uses ARP to determine a physical address for 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 must send a packet to a host but knows only the host IP address, the network station uses ARP to determine the physical address 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 maps the host IP address to its physical address, and then saves the results in an address resolution table for future use.
5. The ARP table on the network station displays the association of the known MAC addresses to IP addresses.

You can configure the lifetime for the learned MAC addresses. The switch executes ARP lookups after 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 static ARP entries. 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.

The switch can use Static ARP entries to solve the following instances encountered on many networks:

- To communicate with a device that does not respond to an ARP request
- To prevent an existing ARP entry from aging out

When you configure a static ARP entry, both the IP address and MAC address of a device is assigned to a physical port. This includes the VLAN number if the physical port is associated with a VLAN.

Proxy ARP

The Ethernet Routing Switch uses Proxy ARP to respond to an ARP request from a locally attached host that is intended for a remote destination. The switch sends an ARP response back to the local host with the MAC address of the switch interface that connects to the host subnet. The switch generates the reply only if the switch has an active route to the destination network.

With you enable Proxy ARP, 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 5000 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 5000 Series intercepts the message, and then responds to the ARP request with the IP address for host C but with the MAC address of the switch itself. Host B then updates its ARP table with the received information.

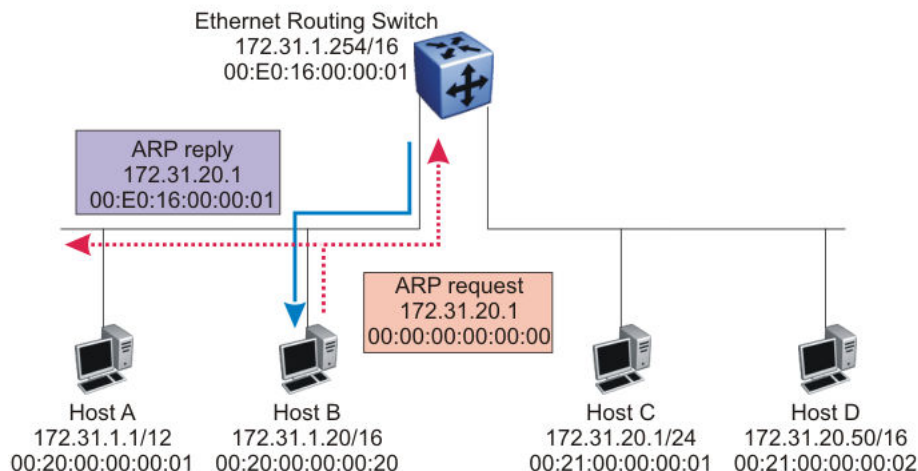


Figure 16: Proxy ARP Operation

Use 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. Do not use Proxy ARP 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

Along with IP Routing, you can use Blocking Mode in two modes: full and none. The following paragraphs explain how blocking mode acts for a stack.

You have a stack with IP Routing enabled and some Layer 3 VLANs. Assign VLANs ports from all the units. Configure IP blocking-mode to Full on the base unit. Remove all the units from stack. All of the units will run in Layer 2 mode. No Layer 3 settings will be available on these units.

You have a stack with IP Routing enabled, and some Layer 3 VLANs. Assign VLAN ports from all the units. Configure the IP blocking-mode to None on the base unit. Remove all of the units from stack. The Layer 3 settings made on the stack will be available on these units. By default IP blocking-mode is None.

IP blocking for stacks

IP Blocking is a Layer 3 feature of the Avaya Ethernet Routing Switch 5000 Series that provides safeguards for a stack where Layer 3 VLANs have port members across multiple stack units. The switch uses IP Blocking whenever a unit leaves a stack or is rebooting inside the context of a stack. Depending on the configuration, Layer 3 functionality is either continued or blocked by this feature.

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

If you configure IP blocking to full, if any units leave the stack, those units run in Layer 2 mode. No Layer 3 settings remain on the units.

If you configure IP blocking 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. After this change 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 can occur to routing traffic until end stations update their ARP cache with the new router MAC addresses. The Avaya Ethernet Routing Switch 5000 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).

Virtual Router Redundancy Protocol (VRRP)

Because end stations are often configured with a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks.

The Virtual Router Redundancy Protocol (VRRP) (RFC 2338) eliminates the single point of failure that can occur after the single static default gateway router for an end station is lost. VRRP introduces the concept of a virtual IP address (transparent to users) shared between two or more routers that connect a common subnet to the enterprise network. Because end hosts use 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, an 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. After 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

When 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. After you enable additional routers in the VLAN, an election process takes place among them to elect a master, 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 virtual MAC address on the Avaya Ethernet Routing Switch 5000 Series is automatically configured as:

```
00-00-5E-00-01-<VRID>
```

where <VRID> is a hexadecimal value in the range 1 to FF that represents the virtual router identification.

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. The backup router does not respond to ARP requests and must discard packets with a MAC address equal to the virtual router MAC address. The backup router does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, the router 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. The backup router sends the VRRP advertisement and ARP request as described in the preceding paragraphs, and then transitions to the controlling state. The virtual router IP address and MAC address does not change, thereby providing transparent operation

VRRP topology example

The following figure shows a VRRP topology example.

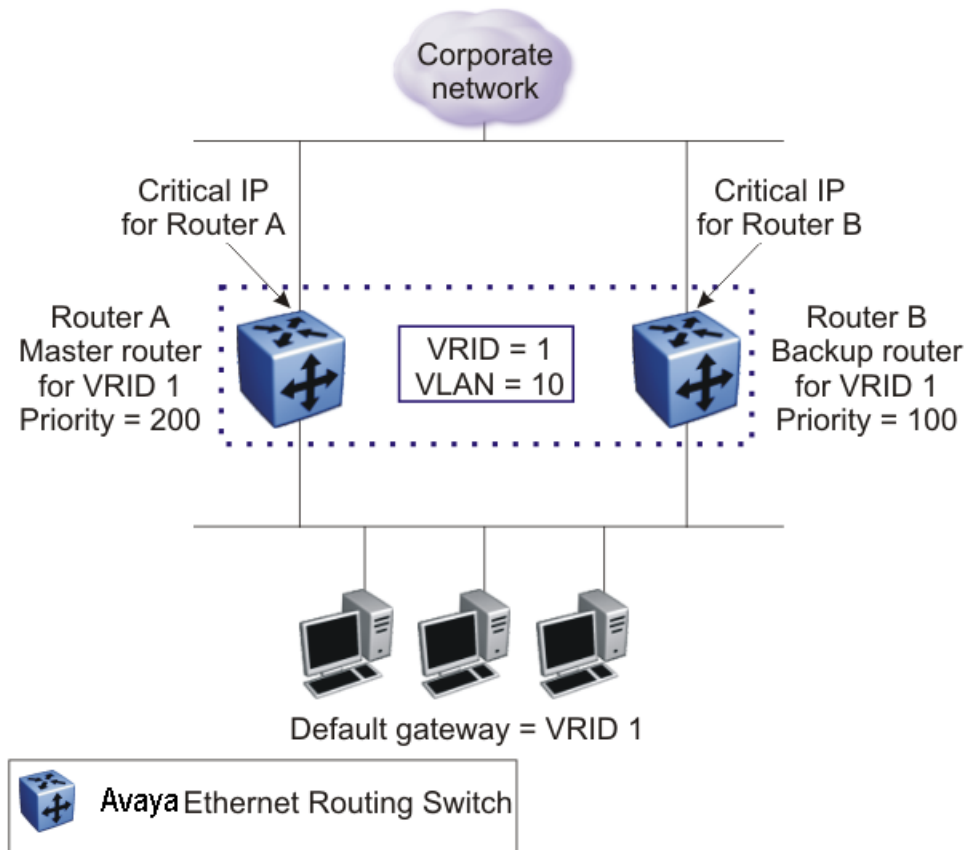


Figure 17: VRRP topology example

In this example, to configure router A as the master router and router B as the backup router, you can configure them for VRRP 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.
 - Configure the virtual router IP address to a previously unassigned IP address.
 - Configure the priority to a value above the priority of 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.
 - Configure the priority to a value below that on Router A, in this case, 100.

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

Critical IP address

Within a VRRP VLAN, it is possible for one link to become inactive, 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 that connects 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 [Figure 17: VRRP topology example](#) on page 67, 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 and SMLT

The standard implementation of VRRP allows only one active master switch for each IP subnet. All other VRRP interfaces in a network are in backup mode.

However, a deficiency occurs when VRRP-enabled switches use Split Multi-Link Trunking (SMLT).

Normally, if a switch connects to two SMLT aggregation switches, the end host traffic is load-shared on all uplinks to the aggregation switches, based on the MLT traffic distribution algorithm. However, VRRP can only have one active routing interface enabled. All other VRRP routers are in backup mode. Therefore, if the SMLT aggregation switches run VRRP, all traffic that reaches the backup VRRP router is forwarded over the Inter Switch Trunking (IST) link

towards the master VRRP router. In this case, the IST link might not have enough bandwidth to carry all the aggregated traffic.

You can overcome this issue by assigning the backup router as the backup master router. The backup master router is a backup router that can actively load-share the routing traffic with a master router.

After you enable the backup master router, the incoming host traffic can be load-shared over the SMLT links as normal. This configuration allows both switches to respond to ARP requests and forward traffic

The following figure shows a sample VRRP configuration with SMLT. Router B is the backup master and the two devices load-share the traffic routing.

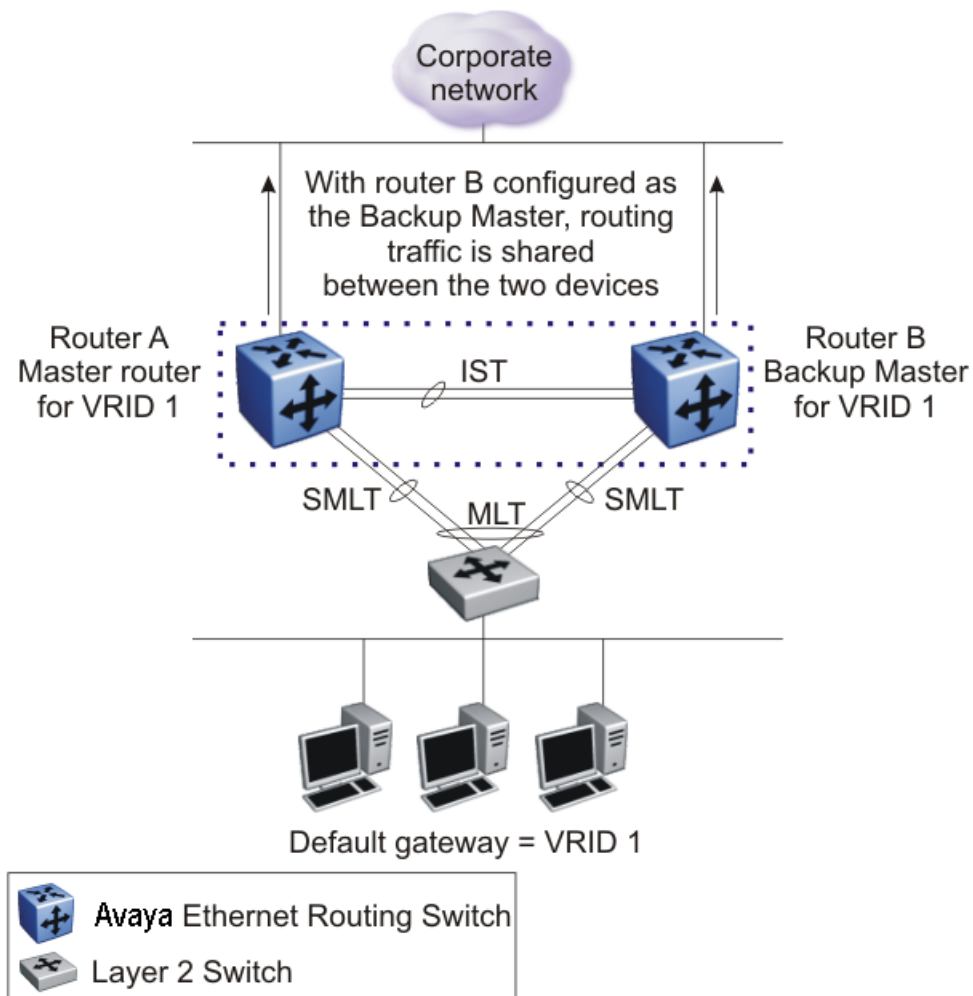


Figure 18: VRRP with SMLT

VRRP fast advertisement interval

With VRRP, you can configure the advertisement interval between sending advertisement messages. Specify the interval value in seconds. This interval permits faster network convergence with standardized VRRP failover. However, losing connections to servers for more than one 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, Avaya Ethernet Routing Switch 5000 Series 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.

After you enable the fast advertisement interval, VRRP can only communicate with other Ethernet Routing Switch devices with the same configuration.

VRF Lite

Virtual Routing and Forwarding (VRF) allows multiple instances of a routing table to co-exist in the same router at the same time. Because the routing instances are independent, the same, or overlapping IP addresses can be used without conflicting with each other.

You can use VRF Lite to maintain networking capabilities and traffic isolation for clients that operate over the same node or router. With VRF Lite, you can create virtual routers to perform the functions of many routers using a single platform. Each virtual router emulates a dedicated hardware router on your network, by providing separate routing functionality to network clients. The result is a substantial reduction in the cost associated with providing routing and traffic isolation for multiple clients.


*** Note:**

You must install the Premier license to enable VRF.

! Important:

The VRF Lite feature is supported on the ERS 5600 series only.

The following figure shows one platform that acts as multiple virtual routers, each serving a different client network.



One ERS 5600 can support up to four virtual routers, with each virtual router instance termed a VRF instance. The switch maintains a separate routing table for each VRF instance.

The switch automatically creates the default, or global VRF instance (VRF 0) at startup. VRF 0 provides all non-virtual and traditional routing services. You cannot delete VRF 0.

You can associate each VRF instance with multiple VLANs and brouter ports. Each VLAN and brouter port can only be associated with a single VRF instance.

VRF Lite characteristics

Consider the following feature characteristics when you configure VRF Lite on the ERS 5600 series:

- Only VRF 0 supports dynamic routing.
- All VRF instances support static routing.
- All VRF instances support DHCP relay.
- Only VRF 0 supports stacking.
- The switch does not support routing between VRF instances.
- VRF instances 1, 2, and 3 are not supported over SMLT.
- Each VRF instance requires a separate uplink, or a tagged uplink.

! Important:

Although VRF instances in an individual switch support duplicate IP addresses, you must ensure that duplicate IP addresses do not exist between VRF instances across switches in the same network.

Circuitless IP

Circuitless IP (CLIP) is a virtual IP (VIP), or loopback interface that provides a method to assign one or more IP addresses to a routing switch, without the requirement of binding the IP address to a physical interface.

Because the IP address assigned to a CLIP interface does not map to a specific physical interface, if one or more physical IP interfaces on a routing switch fails, the CLIP interface ensures connectivity if an actual path is available to reach the device.

The system treats a CLIP interface the same as any IP interface. The network associated with a CLIP is treated as a locally-connected network to the switch, and is always reachable through a VLAN interface. This route always exists and the circuit is always available because there is no physical attachment.

Note:

CLIP interfaces are disabled by default on ERS 5000 Series devices.

CLIP supports the following applications and protocols:

- Internet Control Message Protocol (ICMP)
- Telnet
- Simple Network Management Protocol (SNMP)
- Remote Monitoring (RMON)
- Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP)

You can use a CLIP address as the source IP address in the IP header to send RMON traps.

The system also advertises loopback routes to other routers in the domain, either as external routes using the route-redistribution process, or after you enable OSPF in passive mode, to advertise an OSPF internal route.

CLIP feature considerations

Before you configure CLIP interfaces in your network, consider the following:

- For CLIP interfaces to function properly, you must enable IP routing globally.
- In a stack environment, you can only configure CLIP by using a connection to the base unit.
- Each ERS 5000 Series device supports a maximum of 16 CLIP interfaces.
- CLIP interfaces do not support multinetting.

- A network associated with a CLIP cannot route data traffic.
- RIP does not function on CLIP interfaces, but you can configure RIP routing policies to redistribute CLIP network information.
- OSPF configured on a CLIP interface always runs in passive mode.
- ARP does not function on CLIP interfaces.
- CLIP interfaces do not support Protocol Independent Multicast, Sparse Mode (PIM-SM) or Source Specific Mode (PIM-SSM) .

Chapter 4: IP multicast fundamentals

To manage multicast traffic, the Ethernet Routing Switch 5000 Series supports PIM-SM (for IGMPv1 and IGMPv2), PIM-SSM (for IGMPv3) and IGMP snooping (for IGMPv1, IGMPv2, and IGMPv3). You can enable IGMP snooping on a per-VLAN basis either on a Layer 2 or a Layer 3 VLAN. You can enable PIM-SM and PIM SSM on Layer 3 VLANs only.

This chapter describes the fundamentals of IP multicast as they apply to the Ethernet Routing Switch 5000 Series.

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 1Mbit/sec live video stream for each client, a 100Mbit/sec NIC card on the server can be completely saturated after 90 client connections. The following figure shows an example of this waste of resources.




Figure 19: 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.

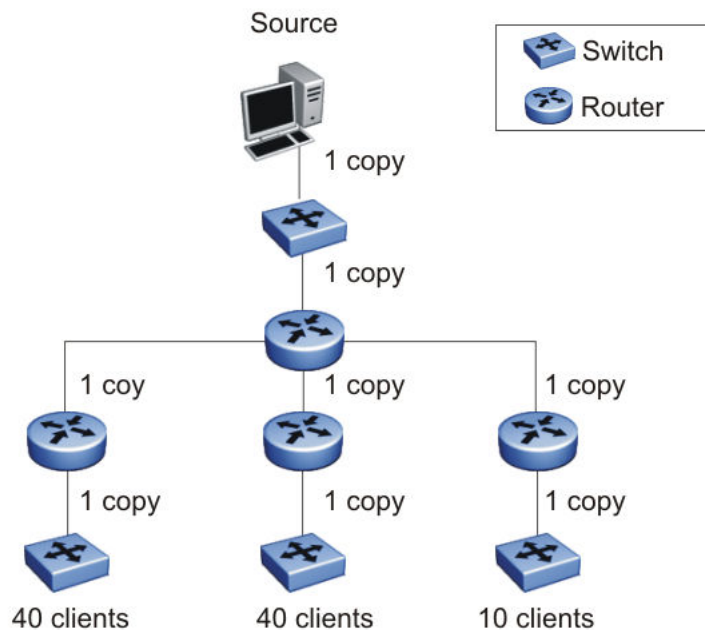


Figure 20: One stream replicated using multicasting

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 UDP packets, which are not as reliable as TCP packets.

Applications that use multicasting to transmit data include:

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

Related topics:

[Multicast groups](#) on page 78

[Multicast distribution trees](#) on page 78

[Multicast addresses](#) on page 80

[IP multicast address ranges](#) on page 80

[IP to Ethernet multicast MAC mapping](#) on page 81

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.

When a multicast enabled router receives a request from a client to receive multicast traffic for a specific group, the system creates an entry of type (*, G) in its mroute table. This indicates that a client is interested in receiving traffic from any source for group G. When the source starts to transmit multicast traffic, an entry of type (S, G) is created. For PIM-SSM only entries of type (S, G) are created.

Multicast distribution trees

A multicast distribution tree consists of the routers that forward multicast data from a particular source to all registered multicast group members. When all routers in the multicast network have determined their upstream and downstream interfaces for a particular group, the multicast tree is formed. At the root of the tree is the source of the multicast stream, and the branches are the destination routers that want to receive the multicast stream. Different multicast protocols use different techniques to discover delivery paths.

The following figure is an example of a simple distribution tree, where the arrows indicate the multicast delivery path from the source to the group members along the multicast distribution tree.




Figure 21: Multicast distribution tree

Related topics:

[Reverse Path Forwarding](#) on page 79

Reverse Path Forwarding

Reverse Path Forwarding is the means by which multicast routers ensure a loop-free topology. When a multicast packet arrives on an interface, the router compares the source address of the packet against the unicast routing table to determine whether the receiving interface is on the shortest path back to the source. If the receiving interface is the one the router would use to forward a unicast packet back to the source, the reverse path check passes, and the router forwards the multicast packet to its downstream neighbors. If the packet does not arrive on an upstream interface, the router discards the packet.

For example, in the preceding figure, if router E receives a packet from source S through router B, which is not on the optimal path back to the source, router E discards the packet. However, if router E receives the source packet from router D, which according to the unicast routing table is the next hop toward the source, router E forwards the incoming packet downstream to router G.

Without reverse path forwarding, loops can form in the network. For example, Router E can forward all packets coming from router D to router B, and router B can in turn forward the same traffic to router D, thereby causing a loop. With the RPF checks running on these routers, no loop can form.

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.

The multicast address range 224.0.0.0/24, 224.128.0.0/24, 225.0.0.0/24, 225.128.0.0/24 up to 239.0.0.0/24, 239.128.0.0/24 maps to reserved multicast MAC addresses. You cannot use these addresses for multicast data traffic.

IP multicast address ranges

IP multicast utilizes D class addresses, which range from 224.0.0.0 to 239.255.255.255. Although subnet masks are commonly used to configure IP multicast address ranges, the concept of subnets does not exist for multicast group addresses. Consequently, the usual unicast conventions where you reserve the all 0s subnets, all 1s subnets, all 0s host addresses, and all 1s host addresses do not apply when dealing with the IP multicast range of addresses.

Addresses from 224.0.0.0 through 224.0.0.255 are reserved by the Internet Assigned Numbers Authority (IANA) for link-local network applications. Packets with an address in this range are not forwarded by multicast capable routers by design. For example, Open Shortest Path First (OSPF) uses both 224.0.0.5 and 224.0.0.6 and Virtual Router Redundancy Protocol (VRRP) uses 224.0.0.18 to communicate across a local broadcast network segment. Multicast address 224.0.0.13 is reserved for establishing adjacency between PIM neighbors

IANA has also reserved the range of 224.0.1.0 through 224.0.1.255 for well-known applications. These addresses are also assigned by IANA to specific network applications. For example, the Network Time Protocol (NTP) uses 224.0.1.1 and Mtrace uses 224.0.1.32. RFC 1700 contains a complete list of these reserved numbers.

Multicast addresses in the 232.0.0.0/8 (232.0.0.0 to 232.255.255.255) range are reserved only for source-specific multicast (SSM) applications, such as one-to-many applications. (For more information, see RFC 4607). While this range is the publicly reserved range for SSM applications, private networks can use other address ranges for SSM.

Finally, addresses in the range 239.0.0.0/8 (239.0.0.0 to 239.255.255.255) are administratively scoped addresses, meaning they are reserved for use in private domains and cannot be advertised outside that domain. This multicast range is analogous to the 10.0.0.0/8, 172.16.0.0/20, and 192.168.0.0/16 private address ranges in the unicast IP space.

Technically, a private network can only assign multicast addresses from 224.0.2.0 through 238.255.255.255 to applications that are publicly accessible on the Internet. Multicast

applications that are not publicly accessible can be assigned addresses in the 239.0.0.0/8 range.

IP to Ethernet multicast MAC mapping

Like IP, Ethernet has a range of MAC addresses that natively support Layer 2 multicast capabilities. However, while IP has a total of 28 addressing bits available for multicast addresses, Ethernet has only 23 addressing bits assigned to IP multicast. The multicast MAC address space for Ethernet is much larger than 23 bits, but only a subrange of that larger space is allocated to IP multicast by the Institute of Electrical and Electronics Engineers (IEEE). Because of this difference, 32 IP multicast addresses map to one Ethernet multicast MAC address.

IP multicast addresses map to Ethernet multicast MAC addresses by placing the low-order 23 bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01:00:5E:00:00:00. Thus, more than one multicast address maps to the same Ethernet address. For example, all 32 addresses from 224.1.1.1, 224.129.1.1, 225.1.1.1, 225.129.1.1, and so on up to 239.1.1.1 and 239.129.1.1 map to the same 01:00:5E:01:01:01 multicast MAC address.

The following figure shows the mapping of multicast IP addresses to MAC addresses.

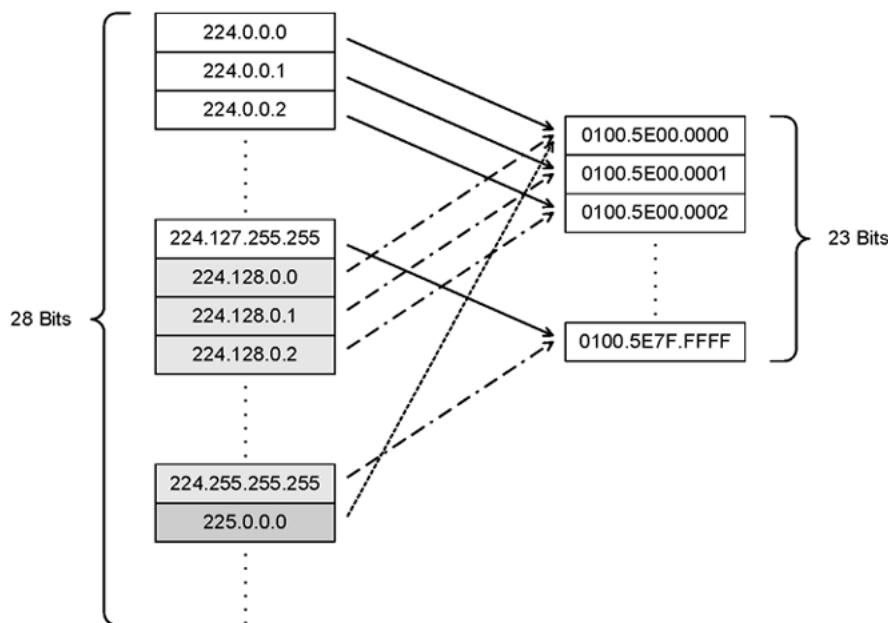


Figure 22: Multicast IP address to MAC address mapping

Most pure Layer 2 Ethernet switches handle Ethernet multicast by mapping a multicast MAC address to multiple switch ports in the MAC address table. However, the Ethernet Routing Switch 5000 Series switches IP multicast data based on the IP multicast address and not the MAC address. It internally maps IP multicast group addresses to the ports that contain group members. After an IP multicast packet is received, the lookup is based on IP group address,

regardless of whether the VLAN is bridged or routed. This avoids the ambiguity in mapping 32 IP addresses to one MAC address.

If your network includes pure Layer 2 Ethernet switches that map each multicast MAC address to 32 IP addresses, the easiest way to avoid any potential issues is to use only a consecutive range of IP multicast addresses corresponding to the lower order 23 bits of that range. For example, use an address range from 239.0.2.0 through 239.127.255.255. A group address range of this size can accommodate the addressing needs of even the largest private enterprise.

Internet Group Management Protocol

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

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. If 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.

From Release 6.2 onwards ERS 5000 Series switches support a maximum of 256 interfaces that can have IGMP associated with them. You can configure IGMP association with or dissociation from an interface. For more information, see [IGMP snooping configuration using ACLI](#) on page 427

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

Related topics:

[IGMPv1 operation](#) on page 83

[IGMPv2 operation](#) on page 84

[IGMPv3 operation](#) on page 87

[IGMPv3 membership report](#) on page 88

[IGMPv3 membership query](#) on page 89

[Multicast flow over Multi-Link Trunking](#) on page 90

[IGMP requests for comment](#) on page 90

IGMPv1 operation

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

IGMPv1 supports two different 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 the 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 to send 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.

Related topics:

[IGMPv1 leave process](#) on page 83

[Host report suppression](#) on page 84

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 participation. If 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 left on the subnet.

Host report suppression

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




Figure 23: IGMP report suppression

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 IGMPv1.

IGMPv2 also implements a querier election process.

IGMPv2 adds support for 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 multicast address: 224.0.0.2)

IGMPv2 also supports IGMPv1 messages.

Related topics:

[Host leave process](#) on page 85

[Maximum Response Time](#) on page 85

[Robustness value](#) on page 86

[Router alert](#) on page 86

[Querier election process](#) on page 86

Host leave process

With IGMPv2, if the host that issued the most recent report leaves a group, it 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.

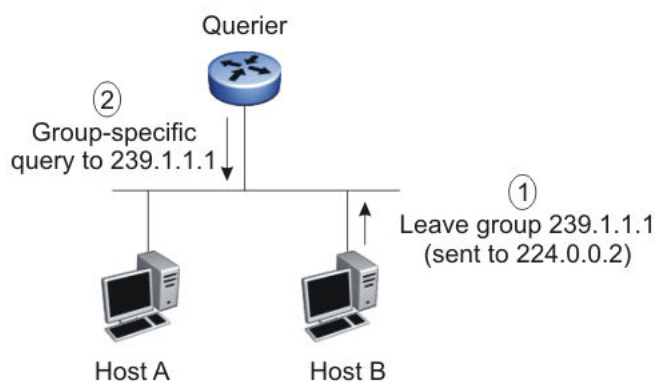


Figure 24: IGMPv2

In this example:

- 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

Maximum Response Time

Each IGMPv2 query from a router to a host includes a Maximum Response Time field, specifying the maximum time n in tenths of a second within which the host must issue a reply.

The host uses this value to calculate a random value between 0 and n tenths of a second for the period that it waits before sending a response.

IGMPv1 queries do not specify a maximum response time. Instead, the maximum response time is a fixed value of 100, that is, 10 seconds.

Robustness value

As part of IGMP configuration, the robustness value lets you configure the switch to offset expected packet loss on a subnet. If you expect a network to lose IGMP query or membership report packets, you can increase the robustness value to offset the lost packets.

When the Ethernet Routing Switch receives an IGMP report from a host, the switch refreshes the expiration time for the group member. The timeout for a group member is a function of the query interval, robustness value, and the maximum response time. If the robustness value is increased, the group member lifetime increases as well, according to the following formula:

IGMP group member lifetime = (IGMP query interval * robustness value) + maximum response time.

If the network is congested, the switch is more likely to miss an IGMP report from a host, or the host can miss the query from the router. To offset the network loss, you can increase the robustness value to extend the life time for the group members.

Router alert

The router alert feature instructs the router to drop control packets that do not have the router-alert flag in the IP header. You can use this option to optimize the performance of multicast routers. When you enable router alert, a router needs to examine only the packets that have the router-alert option flagged, and can ignore all other multicast control packets destined to it. This optimizes the performance in packet processing. This feature is especially useful in routers-only networks. (It is difficult to force a host to send packets with the router-alert option.)

Querier election process

There is normally only one querier 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 hears a query message from a router with a lower IP address, it becomes a nonquerier on that network.

IGMPv3 operation

IGMPv3 adds support for source filtering. The IGMPv3 host can report its interest in receiving multicast packets from only specific source addresses, or the host can report its interest in receiving multicast packets from all but specific source addresses.

IGMPv3 is mostly used in voice and video conferences where multiple people can be part of the same conference. The IGMPv3 packet format adds a v3 Report message type (0x22) and also includes Source-and-Group-specific Query messages.

The message type for Source-and-Group-specific Query message is 0x11, the same as IGMPv1 and IGMPv2. The different Query message versions are identified as follows:

- If the size of the IGMP message type is 8, then it is a v1 or v2 Query message.
- If the Group Address field is 0, then it is a General Query.
- If the Group Address field is a valid multicast IP address, then it is a Group-specific Query.
- If the Group Address field is a valid address and the Number of Sources field is nonzero, then it is a Group-and-Source specific Query message.

Each IGMPv3 Report contains a list of group records. The Group Record contains the multicast group address and the list of source addresses. The record type field specifies whether to INCLUDE or EXCLUDE the list of source addresses that are provided in the Source Address field. For example, to include packets from source 10.10.10.1, the report contains an INCLUDE(10.10.10.1) record.

The list of source addresses can be empty, which is represented by braces ({}), which means either to INCLUDE or EXCLUDE none. For example, the host that wants to receive packets from all group members can send a report with an EXCLUDE({}) record and a host that wants to leave a group can send a report with an INCLUDE({}) record, which is similar to a leave message.

In the following figure, hosts A, B, C, D, E, and F are part of a conference group G1. All hosts except F send a report for group G1 with the mode as INCLUDE(A, B, C, D, E, F) containing all the source addresses. Host F, which is not interested in listening to C and D, sends a report to group G1 with the mode as EXCLUDE(C, D).

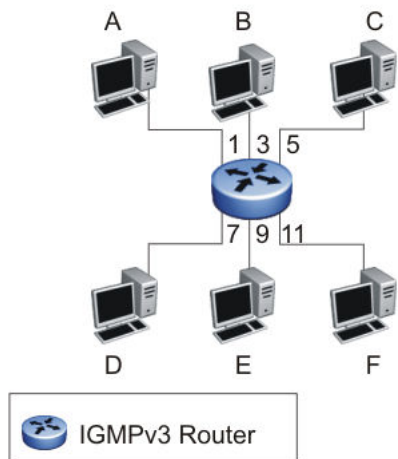


Figure 25: IGMPv3

The router adds the multicast IP address and the list of sources in the forwarding table. The router forwards the packets from A, B, E, and F to all ports. If the packets are received from C and D, it is forwarded to all ports except port 11.

IGMPv3 membership report

IGMPv3 provides the capability to learn which sources are of interest to specific systems, for packets sent to any particular multicast address. IGMPv3 Membership Reports are sent by IP systems to report the current multicast reception state, or changes in the multicast reception state. There are a number of different types of Group Records included in a Report message. The following table shows how IGMPv3 handles the various record types.

IGMPv3 record type	Definition
MODE_IS_INCLUDE (1)	Indicates that the system has a filter mode of INCLUDE for the specified multicast address. The Source Address fields in this Group Record contain the system's source list for the specified multicast address, if it is non-empty.
MODE_IS_EXCLUDE (2)	Indicates that the system has a filter mode of EXCLUDE for the specified multicast address. The Source Address fields in this Group Record contain the system's source list for the specified multicast address, if it is non-empty.
CHANGE_TO_INCLUDE_MODE (3)	Indicates that the system has changed to INCLUDE filter mode for the specified multicast address. The Source Address

IGMPv3 record type	Definition
	fields in this Group Record contain the system's new source list for the specified multicast address, if it is non-empty.
CHANGE_TO_EXCLUDE_MODE (4)	Indicates that the system has changed to EXCLUDE filter mode for the specified multicast address. The Source Address fields in this Group Record contain the system's new source list for the specified multicast address, if it is non-empty.
ALLOW_NEW_SOURCES (5)	Indicates that the Source Address fields in this Group Record contain a list of the additional sources that the system wishes to hear from, for packets sent to the specified multicast address. If the change was to an INCLUDE source list, these are the addresses that were added to the list; if the change was to an EXCLUDE source list, these are the addresses that were deleted from the list.
BLOCK_OLD_SOURCES (6)	Indicates that the Source Address fields in this Group Record contain a list of the sources that the system no longer wishes to hear from, for packets sent to the specified multicast address. If the change was to an INCLUDE source list, these are the addresses that were deleted from the list; if the change was to an EXCLUDE source list, these are the addresses that were added to the list.

IGMPv3 membership query

IP multicast routers send membership queries to query whether hosts are interested in receiving traffic from multicast groups. Specifically for IGMPv3, the router sends a Group-and-Source-Specific Query to learn if any hosts desire reception of packets sent to a specified multicast address, from any of a specified list of sources. In a Group-and-Source-Specific Query, the Group Address field contains the multicast address of interest, and the Source Address fields contain the source addresses of interest. The router can also send Group-Specific Queries upon removal of a source from the multicast group.

Multicast flow over Multi-Link Trunking

In the current release, the Ethernet Routing Switch 5000 Series supports multicast traffic and control packets only on the base link of the MLT.

IGMP requests for comment

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

- For IGMPv1, see RFC 1112.
- For IGMPv2, see RFC 2236.
- For IGMPv3, see RFC 3376.
- 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 Ethernet Routing Switch 5000 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. In this example, 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 26: IP multicast propagation on a LAN without IGMP snooping

To prune ports that are not group members from receiving the group data, the Ethernet Routing Switch 5000 Series supports IGMP snoop for IGMPv1, IGMPv2, and IGMPv3. With IGMP snoop enabled on a VLAN, the switch forwards the multicast group data to only those ports that are members of the group. Using IGMP snoop, VLANs can provide the same benefit as IP Multicast routers, but in the local area.

The Ethernet Routing Switch 5000 Series 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. Using the information gathered from the reports, the switch builds 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, 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 27: 5000 Series switch running IGMP snooping

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

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 5000 Series switch supports IGMPv3 source filtering capability with IGMPv3 Snooping. IGMPv3 Snooping remains backward compatible with IGMPv1 and IGMPv2.

On the Ethernet Routing Switch 5000 Series, IGMPv3 Snooping-enabled interfaces process IGMP reports as follows:

- process all six IGMPv3 group record types
- process all IGMPv3 source information
- backward compatible with IGMPv1/IGMPv2 reports

Due to limitations on the ERS 5510 switch, IGMPv3 source filtering is not implemented in hardware, and multicast traffic entering the system through an ERS 5510 switch is not filtered

based on source. The ERS 5510 switch is only able to add the multicast IP address, but not the source information, to its forwarding table.

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 5000 Series can consolidate these multiple reports 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, for example if the switch adds another multicast group or receives a query since the last report is transmitted upstream, then the switch forwards a new report to the multicast router ports.

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

In the following figure, switches S1 to S4 represent a 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.

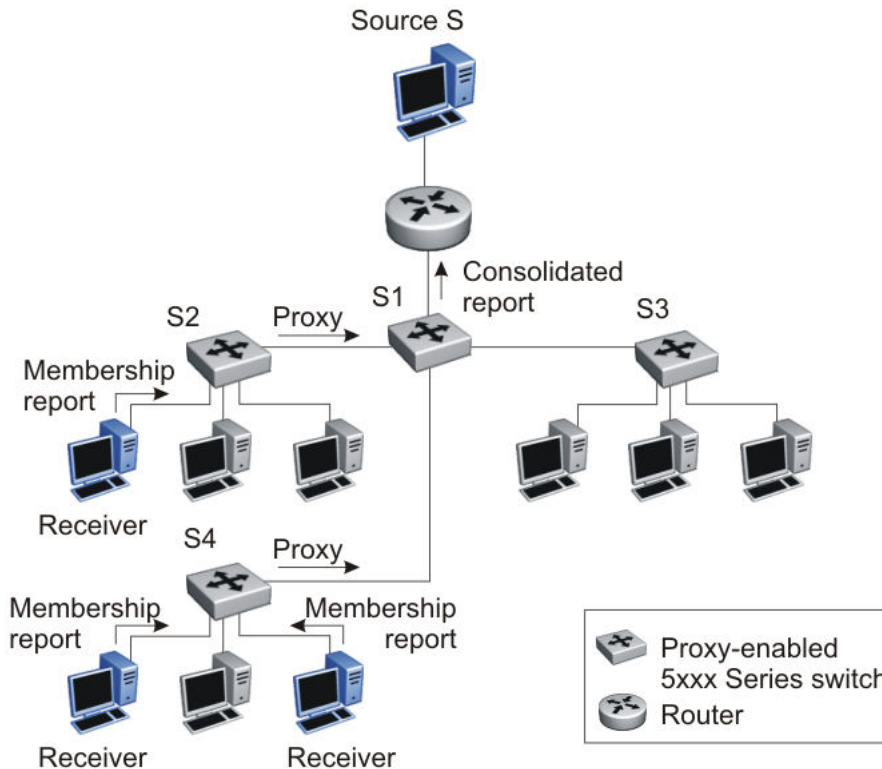


Figure 28: 5000 Series switch running IGMP proxy

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 way, 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 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.

Report forwarding

When forwarding IGMP membership reports from group members, the Ethernet Routing Switch 5000 Series forwards the reports only to those ports where multicast routers are attached. For 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., or to PIM ports.

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.

The following figure shows how static mrouter ports operate. Here, the Ethernet Routing Switch 5000 Series has port members 5/1 and 6/1 connected to IGMP routers in VLAN 10. In this case, router 1 is the IGMP querier because it has a lower IP address than router 2. Router 2 is 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.

! Important:

Configure a static mrouter port only when multiple multicast routers are present that are not directly attached to one another, but are directly attached to the VLAN (technically an invalid configuration). If multicast routers have an existing route between them (the valid configuration) and this field is configured, a multicast loop forms.




Figure 29: Static mrouter port and nonquerier

Unknown multicast packet 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 Ethernet Routing Switch 5000 Series, 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, use the `vlan igmp unknown-mcast-no-flood` CLI command.

With this feature enabled, the switch forwards all unknown multicast traffic to IGMP static and dynamic 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 NVRAM.

Allowing a multicast MAC address to flood VLANs

The unknown multicast filtering feature introduces a potential problem when you place a Layer 2 VLAN between two Layer 3 switches that exchange protocol packets (such as OSPF). Since the protocols do not join a multicast group, the IGMP snooping process cannot identify their multicast MAC addresses. The system drops the protocol packets due to the effect of the unknown multicast filtering feature. The two Layer 3 switches can never establish adjacencies and the OSPF protocol fails.

Using the `vlan igmp unknown-mcast-allow-flood` CLI command, you can specify multicast MAC addresses or multicast IP addresses that need to be flooded on a VLAN even when the unknown multicast filtering feature is enabled. The specified MAC or IP addresses are added to the allow-flood table for the specified VLAN. Any matching packets are flooded on all ports of a VLAN.

Because multicast MAC addresses starting with 01:00:5E map to multiple IP addresses, you cannot specify 01:00:5E MAC addresses in the allow-flood table. Instead, you must specify the required multicast IP address to flood. For instance, you cannot add MAC address 01.00.5E.01.02.03 to the allow-flood table, but you can add IP address 224.1.2.3.

For all other types of MAC address, you can enter the MAC address directly to allow flooding. For example, to allow flooding of STP BPDUs, you can specify MAC address 01:80:c2:00:00:00.

IGMP snooping configuration rules

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

- While PIM-SM supports up to 1000 multicast groups, the maximum number of groups supported with IGMP snooping are as follows:
 - Ethernet Routing Switch 5510: 240 groups

- Ethernet Routing Switch 5520 and 5530: 492 groups
- Ethernet Routing Switch 5600: 992 groups

In a hybrid stack that includes ERS5510 units, you can use the IGMP mode to set ERS5510 capability for the stack. ERS5510 mode supports up to 240 groups. Or you can configure the non-5510 mode.

When you configure the non-5510 mode, the system disables IGMP functionality on the ERS5510, and the system automatically sets the maximum number of groups to 492, if an ERS 5520 or ERS 5530 is present in the stack (hybrid), or to 992 groups for a pure ERS 5600 stack.

The system purges existing groups when you change the IGMP mode.

If the multicast group table reaches its limit, the system cannot add a new entry with a JOIN message or add a new sender identifying a new group and the hardware discards the multicast stream from the new sender. New entries can be added again once the table is not full.

Because all multicast forwarding shares the same hardware resources, when you deploy PIM the actual number of supported IGMP groups can be reduced.

- When you use IGMPv1 or IGMPv2 Snooping mode, the maximum number of IGMP groups is determined by the IGMP operational mode. When you use IGMPv3 Snooping mode, the maximum number of IGMP groups is determined by the types of IGMPv3 report packets processed. Each new IGMPv3 group or source consumes 1 hardware table entry. If the IGMPv3 source is for a group that has never been seen by the IGMP application, then an additional hardware table entry is consumed.
- When you specify MAC addresses or IP addresses to be flooded on the switch, the specified MAC or IP addresses are flooded on all ports of the specified VLAN. In addition, if multicast join messages are received for IP addresses specified in the allow-flood table, these IP addresses are not displayed in the IGMP group membership table. In other words, the switch does not learn groups if they are specified in the allow-flood table.
- A port that is configured for port mirroring cannot be configured as a static mrouter port.

The switch does support mirroring of IGMP control packets as well as multicast data packets.

- If a Multi-Link Trunk member is configured as a static mrouter port, all of the Multi-Link Trunk members are configured as static mrouter ports. Also, if you remove a static mrouter port, and it is a Multi-Link Trunk member, all Multi-Link Trunk members are automatically removed as static mrouter port members.
- Static mrouter ports must be port members of at least one VLAN.
- The IGMP snooping feature is not STP dependent.
- The IGMP snooping feature is not Rate Limiting dependent.
- The snooping feature must be enabled for the proxy feature to have any valid meaning.
- Static mrouter ports are configured per VLAN.

- Mrouter ports cannot be configured on ports configured for LACP (LACP is mutually exclusive with mrouter port).
- The Ethernet Routing Switch 5510 does not support IGMP with Multiple Spanning Tree.

! Important:

Because IGMP snooping is set up per VLAN, all IGMP changes are implemented according to the VLAN configuration for the specified ports.

IGMP and stacking

All IGMP features that are supported in the standalone mode are also supported in stacking mode. The configuration of IGMP from the ACLI is supported only from the base unit. This behavior is similar to all the other Layer 3 functions and routing protocols. However, it is different from the IGMP implementation in previous releases.

Default IGMP values

Parameters	Range	Default Value
Snooping	Enable/Disable	Disable
Version	1–3	2
Proxy	Enable/Disable	Disable
Query Interval	0–65535	125
Last Member Query Interval	0–255	10
Query Max. Response Time	0–255	100
Robust Value	2–255	2
Router Alert	Enable/Disable	Disable
Multicast Router ports	Port masks	Disable port masks
Op-mode	5510/non-5510	5510

IGMP snooping interworking with Windows clients

This section describes an interworking issue between Windows clients and the Ethernet Routing Switch 5000 Series when you enable IGMP snooping for multicast traffic.

Under normal IGMP snooping 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, initially reply with IGMPv2 reports. However, after a period of time, the Windows clients switch to IGMPv3 reports, which the Ethernet Routing Switch 5000 Series does not recognize. In this case, the switch prunes the Windows client from the group and only forwards traffic to any non-Microsoft clients that are left in the group. If no other group members are left, the switch can revert to flooding all ports (in which case, the Windows client still receives the stream). Alternatively, the switch may be pruned altogether from the multicast group (in which case, the Windows client no longer receives the stream.)

To force a Windows client to only use IGMPv1 or IGMPv2 reports so that these symptoms do not occur, 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 utilize IGMPv2 only, 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 it 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.

IGMP Send Query

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.

Avaya recommends that each VLAN using IGMP multicast have a router performing multicast queries. This router typically has PIM or DVMRP enabled. Currently, PIM is only available for standalone devices. Networks with no standalone devices currently have no capability for implementing the pruning of IGMP traffic. The IGMP Send Query 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 Send Query 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 Send Query functionality can only be enabled when IGMP snooping is active on the switch or stack.

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

Protocol Independent Multicast-Sparse Mode

Protocol Independent Multicast-Sparse Mode (PIM-SM), as defined in RFC 2362, supports multicast groups spread out across large areas of a company or the Internet. Unlike dense-mode protocols, such as Distance Vector Multicast Routing Protocol (DVMRP), that initially flood multicast traffic to all routers over an entire internetwork, PIM-SM sends multicast traffic only to routers that belong to a specific multicast group and that choose to receive the traffic. This technique reduces traffic flow over wide area network (WAN) links and minimizes the overhead costs of processing unwanted multicast packets.

Dense-mode protocols that use the flood-and-prune technique are efficient when receivers are densely populated; however, for sparsely populated networks, PIM-SM is more efficient.

PIM-SM is independent of any specific unicast routing protocol, but it does require the presence of a unicast routing protocol, such as RIP or OSPF. PIM-SM uses the information from the unicast routing table to create and maintain multicast trees that allow PIM-enabled routers to communicate.

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

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

Related topics:

[PIM-SM concepts and terminology](#) on page 101

[PIM-SM shared trees and shortest-path trees](#) on page 105

[Source-to-RP SPT](#) on page 108

[Register suppression timeout](#) on page 109

[Receivers leaving a group](#) on page 109

[PIM assert](#) on page 109

PIM-SM concepts and terminology

The following sections describe PIM-SM concepts and terminology.

Related topics:

[PIM-SM sources and receivers](#) on page 101

[PIM neighbor discovery](#) on page 101

[Required elements for PIM-SM operation](#) on page 102

[Designated router](#) on page 102

[Rendezvous-point router](#) on page 103

[Active RP selection](#) on page 103

[Static RP](#) on page 104

[Bootstrap router](#) on page 105

[Active BSR selection](#) on page 105

PIM-SM sources and receivers

With PIM-SM, a host can be a source, a receiver, or both:

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

PIM neighbor discovery

To discover neighbors, PIM routers exchange PIM hello packets. When PIM is enabled on a router interface, the interface forwards PIM hello packets to the all-PIM-Routers multicast address (224.0.0.13).

Each PIM hello packet contains a holdtime that specifies the period that the receiving router must wait before declaring the neighbor unreachable. This holdtime is configurable as the

query interval for each interface. Each PIM interface continues to send hello messages at the configured query interval.

Required elements for PIM-SM operation

PIM-SM operates in a domain of contiguous routers that have PIM-SM enabled. Each router must run an underlying unicast routing protocol to provide routing table information to PIM-SM.

Each PIM-SM domain requires the following routers:

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

Within the PIM-SM domain, each group can have only one active RP router and one active BSR. The active BSR is chosen among a list of candidate-BSRs, and the active RP is chosen among a list of candidate-RPs. You can configure the Ethernet Routing Switch 5000 Series to be a candidate-BSR, a candidate-RP, or both.

Designated router

The designated router (DR) serves as the link from sources and receivers to the other routers in the PIM-SM domain. There are typically multiple DRs in a PIM-SM domain.

On any subnet, the DR is the PIM-SM router with the highest IP address. The DR performs the following tasks:

- sends register messages to the RP router on behalf of directly connected sources
- sends join/prune messages to the upstream router on behalf of directly connected receivers
- maintains information about the status of the active RP router

! Important:

You cannot manually configure a router as a DR. If a router is enabled with PIM-SM and it is the PIM-SM router with the highest IP address on the subnet, it automatically acts as the DR for any directly attached sources and receivers, as required.

Rendezvous-point router

A multicast group has only one active rendezvous-point (RP) router. The RP performs the following tasks:

- manages one or several IP Multicast groups
- becomes the root for the shared tree to these groups
- accepts join messages from receivers
- registers sources that want to send data to group members
- forwards data to the group

At the RP router, receivers meet new sources. Sources register with the RP to identify themselves to other routers on the network; receivers join the RP-based multicast distribution tree to learn about new sources.

For each multicast group, PIM-SM builds a multicast distribution tree, known as the shared tree, with the RP at the root and all receivers downstream from the RP. Although you can physically locate the RP anywhere on the network, the RP must be as close to the source as possible.

Active RP selection

The active RP is calculated among a list of candidate RPs (C-RP). Within each group, you can configure multiple PIM-SM routers as C-RPs.

Each C-RP sends unicast advertisement messages to the BSR. The BSR creates a list of C-RPs, which is referred to as the RP set. The BSR periodically sends bootstrap messages that contain the complete RP set to all routers in the group. Each router uses the same hash function to determine which router in the set is going to be the RP (given the same RP set, each router points to the same RP). If the active RP fails, routers can recalculate the active RP using the reduced set of C-RPs.

You can only configure one RP candidate on a multicast enabled router for a single group or for a range of groups. If you configure multiple RP-candidates for the same group range they are all used as active RPs. The election is made by the BSR using a hash algorithm.

The active BSR sends a list with all the active RP set configured in the PIM domain to all PIM-SM enabled routers. A router that receives a Join request creates a (*, G) group type entry in the mroute table only if an active RP exists for group G.

A router that was elected as active RP for a group will have in the mroute table all entries of type (*, G) and (S, G) for that group.

Static RP

You can use the static RP feature to configure a static entry for an RP. Static RP-enabled routers do not learn about C-RPs through the BSR. With static RP enabled, the router ignores BSR messages and loses all dynamically learned BSR information. When you configure static RP entries, the router adds them to the RP set as though they are learned through the BSR.

You can use the static RP feature when dynamic learning is not needed, typically in small networks or for security reasons. You can also enable static RP to allow communication with routers from other vendors that do not use the BSR mechanism. Some vendors use early implementations of PIM-SMv1 that do not support the BSR or proprietary mechanisms like the Cisco Auto-RP. For a network to work properly with static RP, all the routers in the network (including routers from other vendors) must be configured with the same RP or RPs, if several RPs are present in the network.

To configure static RP on a router, the next hop of the unicast route toward the static RP must be a PIM-SM neighbor. If a route change causes the next hop toward an already configured static RP to become a non-PIM neighbor, the PIM-SM protocol fails on the router. The state of the configured RP on the router remains invalid until it can be reached through a PIM neighbor.

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

When you configure a static RP, take into account the following considerations:

- You cannot configure a static RP-enabled router as a BSR or as a C-RP.
- All dynamically learned BSR information is lost. However, if you disable static RP, the router clears the static RP information and regains the BSR functionality.
- Static RPs do not age; that is, they cannot time out.
- Routers do not advertise static RPs; therefore, if a new PIM-SM neighbor joins the network, this new neighbor does not know about the static RP unless you configure the neighbor with that static RP.
- All the routers in the network (including routers from other vendors) must map to the same RP.
- In a PIM-SM domain with both static and dynamic RP routers, you cannot configure one of the (local) interfaces of the static RP routers as RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If a mix of Avaya and other vendor routers exist across the network, ensure that all routers use the same active RP because other vendors can use different algorithms to elect the active RP. The Ethernet Routing Switch 5000 Series uses the hash function defined in the PIM-SM standard to elect the active RP, with the highest C-RP address selected to break a tie. Other vendors can use the lowest IP address to break the tie.

- You cannot assign a priority to static RP entries, although the Ethernet Routing Switch accepts priority values from non-Avaya routers for interoperability.
- A static RP that you configure on the router is alive as long as the router has a unicast route to the network for the static RP. If the router loses this route, it invalidates the static RP and uses the hash algorithm to remap all affected groups. If the router regains this route, it validates the static RP and uses the hash algorithm to remap the affected groups.

Bootstrap router

The bootstrap router (BSR) receives advertisement messages from the C-RPs. The BSR adds the C-RPs and their group prefixes to the RP set. The BSR sends bootstrap messages that contain the complete RP set to all routers in the domain to allow them to learn group-to-RP mappings.

Only one BSR exists for each PIM-SM domain.

Active BSR selection

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

PIM-SM shared trees and shortest-path trees

PIM-SM uses two types of multicast distribution trees to deliver data packets to group members: shared trees and shortest-path trees (SPT).

Related topics:

[Shared tree](#) on page 105

[Traffic forwarding with the shared tree](#) on page 106

[Shortest path tree](#) on page 107

[Receiver joining a group and receiving data from a source](#) on page 108

Shared tree

The shared tree connects all members of the multicast group to a central core router, the active RP, which is at the root of the shared tree.

The construction of the shared tree begins when a host sends an IGMP membership report to a local DR to join a multicast group. The DR in turn signals join messages toward the RP. The intermediate routers toward the RP add the group entry when forwarding the join messages. When the join messages reach the RP, the RP adds the tree branch to the shared tree for the group.

Although a shared tree is less efficient than a source-rooted tree, PIM-SM shared tree reduces the network bandwidth during tree construction and maintenance, as flood-and-prune messages are not required.

The following figure shows an example of an RP-based shared tree.

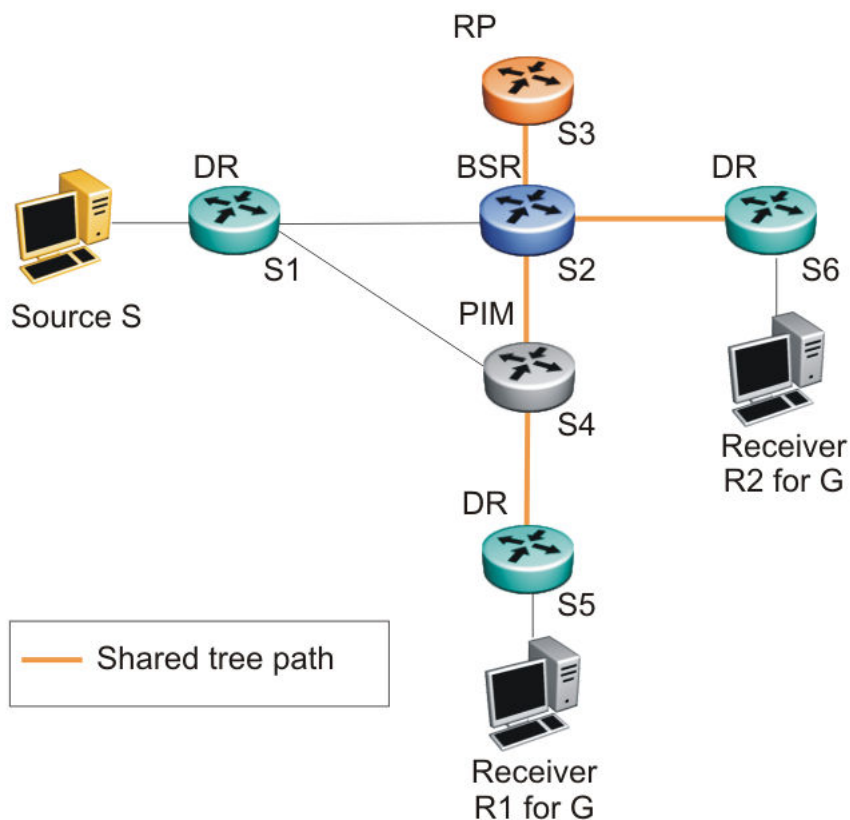


Figure 30: RP-based shared tree

Traffic forwarding with the shared tree

All group traffic initially flows from the RP downstream through the shared tree to the receivers. To forward multicast data from a source to group members, the source DR encapsulates the multicast packets in Register messages that it then unicasts to the RP. The RP decapsulates the Register messages, and then forwards the multicast data to any existing group members downstream using the shared tree.

In the shared tree, the RP router represents a potential bottleneck and a single point of failure. As a result, PIM-SM allows local DRs to bypass the share tree and switch to a source-rooted shortest path tree.

Shortest path tree

When multicast packets arrive at the receiver DR, the DR can identify the IP address of the source. If the DR determines that the shared tree is not the optimal path back to the source, it sends a join message directly to the source DR. This new direct path from the source to the receiver DR is the source-based shortest-path tree (SPT). When the receiver DR starts receiving traffic directly from the source, it sends a prune message to the RP to stop sending messages over the shared tree.

With the Avaya Ethernet Routing Switch 5000 Series switches, the DR switches to the SPT after it receives the first packet from the RP.

The following figure shows an example of a source-based SPT.

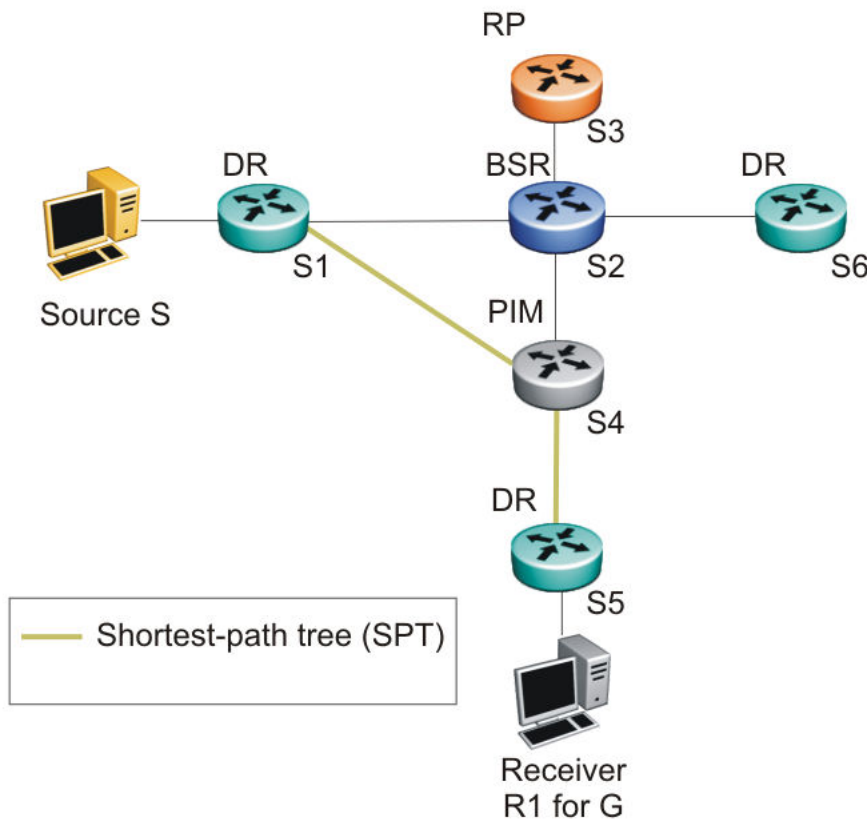


Figure 31: Source-based SPT

Receiver joining a group and receiving data from a source

The following steps describe how the receiver R1 in [Figure 30: RP-based shared tree](#) on page 106 and [Figure 31: Source-based SPT](#) on page 107 joins multicast group G:

1. The BSR distributes RP information to all switches in the network. In this example, based on the RP hash function, S3 is the RP for group G.
2. Receiver R1 multicasts an IGMP host membership report for group G, which the DR (S5) receives.
3. Acting on this report, S5 creates a (*,G) route entry in the multicast forwarding table and sends a (*,G) join to the RP.
4. The intermediate routers toward the RP (S4 and S2) add the (*,G) route entry when forwarding the join message to the RP.
5. The RP adds the port that receives the join as a downstream port for the (*,G) group.
6. The source S starts multicasting data to group G.
7. The source DR (S1) encapsulates the data in a Register message that it unicasts to the RP (S3).
8. S3 decapsulates the multicast data and forwards it down the shared tree. Group member S5 receives the data and forwards it to receiver R1.
9. After S5 receives the first packet, it knows the IP address for the source. S5 creates an (S,G) entry in the multicast forwarding table, and sends a (S,G) join to the source. All intermediate routers along the path to the source create the (S,G) entry. S5 also prunes itself from the RP shared tree.
10. S1 forwards multicast packets to S5 over the SPT.

! Important:

The PIM-SM topology shown in this example is simplified and is not the best design for a network if the source and receiver are placed as shown. In general, RPs are placed as close as possible to sources.

Source-to-RP SPT

Rather than continue to receive multicast traffic from the source through unicast Register messages, the RP also switches to a source-based SPT. After it receives the first source Register message, it sends a join message to the source DR to receive the data through a multicast rather than unicast stream. After it receives the first multicast packet over the SPT, the RP sends a register-stop message to the source to stop sending the data in register messages.

On the Ethernet Routing Switch 5000 Series, the DR only forwards the first multicast packet as a Register packet to the RP, and immediately goes into discard mode until it receives a join message from the RP. During this time, there is brief data loss of the multicast stream.

After the source DR processes the join message, the DR forwards native multicast packets to the RP over the SPT path.

Register suppression timeout

If a source registers with an RP, but no receivers are registered to receive the traffic, the RP sends a register-stop to the source.

After receiving a register-stop message from the RP, the source DR starts a register suppression timer (the default value is 60 seconds).

Shortly before the register suppression timer expires, the source DR sends a register message with no encapsulated packets to the RP router. This null-register message prompts the RP router to determine whether new downstream receivers joined the group. If no new members have joined the group, the RP router sends another register-stop message to the DR for the source, and the register suppression timer restarts. In this way, the DR can regularly poll the RP to determine whether any new members have joined the group without forwarding larger traffic packets to the RP unnecessarily.

A lower register suppression timeout produces traffic bursts from the DR more frequently, whereas with a higher value, new receivers face a longer join latency.

Receivers leaving a group

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

PIM assert

When a PIM router connects a source to a LAN segment and it detects a second PIM router with a route to the same source on the same segment, the routers exchange Assert messages to determine which router is to forward the multicast stream on the segment. The router that is elected after the change of the Assert messages is known as DR (Designated router) and is the one with the highest IP address.

PIM passive interfaces

You can specify whether you want a PIM interface to be active or passive. The default is active. Active interfaces can transmit and receive PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you have a high number of PIM interfaces and these interfaces are connected to end users, not to other routers.

A PIM passive interface drops any messages of the following type:

- Hello
- Join/Prune
- Register
- Register-Stop

! Important:

A device can send Register and Register-Stop messages to a PIM passive interface, but that interface cannot send out these messages.

- Assert
- Candidate-RP-Advertisement
- Bootstrap

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

The PIM passive interface maintains information through the IGMP protocol about hosts that are related to senders and receivers, but the interface does not maintain information about any PIM neighbors.

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

You can configure a PIM passive interface as a BSR or an RP, although Avaya does not recommend these options.

! Important:

Before you change the state (active or passive) of a PIM interface, disable PIM on that interface. Disabling PIM prevents instability in the PIM operations, especially when neighbors are present or streams are received.

PIM-SM capabilities and limitations

The following list describes the capabilities and limitations of PIM-SM on the Ethernet Routing Switch 5000 Series.

- PIM-SM is not supported on the Ethernet Routing Switch 5510 platform.
- PIM-SM is supported in a stack but PIM has no effect on ERS 5510 units in a stack. Native multicast streams and PIM protocol packets ingress on the ERS 5510 are flooded.
- You cannot allow the PIM-SM shared path tree or SPT to span across any Layer 2 switches. Be sure to implement your topology such that the unicast routes from any DR to the PIM RP and to all multicast sources travel through directly-connected PIM neighbors only. Otherwise, network issues may arise.
- PIM-SM cannot be enabled on brouter ports.
- PIM-SM is not supported over SMLT or IST.
- PIM-SM is not supported on a secondary IP of a Layer 3 VLAN.
- A maximum of 32 PIM-SM active interfaces are supported.
- A maximum of 64 total PIM-SM interfaces are supported.
- You can configure only one Candidate-RP per switch for any number of groups (up to 50 group ranges).
- You can configure static RP for up to 50 groups.
- You can configure every PIM-enabled interface as Candidate-BSR.
- PIM-SM supports forwarding of the multicast stream on ECMP, but traffic balancing is not supported. PIM-SM picks one route for its RPF check and uses it for all streams when joining a source on this route.
- On the Ethernet Routing Switch 5600, up to 992 (S,G) entries are supported.
- On the Ethernet Routing Switch 5500, up to 492 (S,G) entries are supported. On a hybrid stack that contains ERS 5500 and ERS 5600 units, up to 492 (S,G) entries are supported.
- Some Layer 2 IGMP snooping-enabled switches can learn a maximum of 240 groups from clients. However, a PIM router can learn more than 240 groups if it is connected to more than one snooping-enabled switch. In this case, if each Layer 2 switch learns 240 groups, the number of groups the PIM router learns is: $240 * \text{number of Layer 2 switches}$. However, the number of (*,G) entries on the PIM router is limited to 960 for a 5600 switch and 490 for a 5500 switch.
- If a PIM server and IGMP receiver are in the same VLAN, you cannot connect them to the same port. To have a PIM server and IGMP receiver on the same port, the server and receiver must be in different tagged VLANs.

- With static RP, priority is not supported in a pure Avaya-only solution. If the 5000 Series switch is connected to a non-Avaya router that is running static RP, then the 5000 Series switch can learn the priority as advertised by the non-Avaya router.
- Passive interfaces are supported on the edge only (where the port only has connections to either clients or servers). Make sure that any passive interfaces are not in the path of any PIM RPF paths, otherwise the network may not work.

Enabling or disabling routing with IGMP enabled

You cannot enable PIM-SM and IGMP snooping on the same VLAN. As a result, when enabling or disabling routing on a VLAN, the IGMP functionality operates as follows:

- When a VLAN is changed from Layer 3 to Layer 2, IGMP reduces the functionality to snooping.
- When an IGMP snooping-enabled VLAN is changed from Layer 2 to Layer 3, the switch continues to support snooping until you enable PIM-SM.
- When a Layer 3 VLAN has v1 and v2 snooping enabled and you try to enable PIM-SM, the switch displays a warning message to disable snooping before enabling PIM-SM.

Nonsupported IGMP features

The following list describes nonsupported IGMP features on the Ethernet Routing Switch 5000 Series.

- Multicast Router Discovery
- Fast Leave
- Channel Limit—Limit the number of groups a host (port) can join at a time
- IGMP Static Address configuration

Default PIM-SM values

The following table describes the PIM-SM default values.

Parameter	Definition	Range	Default Value
Global PIM-SM status	Indicates the status of PIM-SM on the switch.	Enabled/Disabled	Disabled
PIM mode	Specifies the global PIM mode on the switch.	Sparse mode or SSM mode	Sparse mode
Bootstrap Period	At the elected BSR, this is the interval between originating bootstrap messages.	5–32 757 seconds	60 seconds
C-RP Advertise Timeout	Indicates the frequency with which candidate RPs periodically send C-RP-Adv messages.	5–26 214 seconds	60 seconds
Unicast Route Change Timeout	Specifies how often the routing information that PIM uses is updated from the routing table manager (RTM).	2–65 535 seconds	5 seconds
Join/Prune Interval	Indicates how long the switch waits between sending out join/prune messages to the upstream neighbors.	1–18 724 seconds	60 seconds
Register Suppress Timeout	Specifies how often the source DR polls the RP using data packets encapsulated in Register messages.	6–65 535 seconds	60 seconds
Data Discard Timer	After the router forwards the first source packet to the RP, this value specifies how long (in seconds) the router discards subsequent source data while waiting for a join from the RP.	5–65 535 seconds	60 seconds
Static RP	Indicates the status of static RP on the switch.	Enabled/Disabled	Disabled
Forward Cache Timeout	Indicates the PIM-SM forward cache expiry value. This value is used in aging PIM-SM mroutes.	10–86 400 seconds	210 seconds
VLAN PIM-SM status	Indicates the status of PIM-SM on the VLAN.	Enabled/Disabled	Disabled
Hello Interval	Sets the hello interval for the VLAN.	0–18 724 seconds	30 seconds

Parameter	Definition	Range	Default Value
Interface Type	Sets the interface type on a particular VLAN.	<ul style="list-style-type: none"> • active: allows PIM-SM control traffic to be transmitted and received. • passive: prevents PIM-SM control traffic from being transmitted or received. 	Active
Candidate-BSR priority	Indicates whether the router is acting as a C-BSR on a particular VLAN, and if so, the priority associated with it.	0 to 255	-1 (indicates that the interface is not a Candidate BSR)
Candidate-RP	Indicates whether the VLAN interface is configured as a C-RP. With the Ethernet Routing Switch 5000 Series, you can configure only one local interface as a C-RP for any number of groups.	IP address of the C-RP interface and the associated group and mask.	None defined (disabled)

PIM-SSM overview

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

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

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

SSM applications use IP addresses reserved by the Internet Assigned Numbers Authority (IANA) in the 232/8 range (232.0.0.0 to 232.255.255.255). SSM recognizes packets in this

range and controls the behavior of multicast routing devices and hosts that use these addresses. When a source S transmits IP datagrams to an SSM destination address G, a receiver can receive these datagrams by subscribing to the (S,G) channel. A channel is a source-group (S,G) pair where S is the source sending to the multicast group and G is an SSM group address.

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

Avaya recommends running PIM-SSM on either all the switches in the domain or only on the edge routers. If there is a mix of PIM-SSM and PIM-SM switches in the domain, run PIM-SSM on all the edge routers and PIM-SM on all the core routers. A PIM domain with edge routers running PIM-SM and core routers running PIM-SSM does not work properly. SSM switches running IGMPv3 drop any reports that they receive out of the SSM range. The SSM switch does not forward them to a PIM-SM switch.

Related topics:

[PIM SSM concepts and terminology](#) on page 115

[Theory of operation](#) on page 116

PIM SSM concepts and terminology

The default PIM mode is PIM-SM. You can change the mode from PIM-SM to PIM-SSM at any time, however, you can change from SSM to SM only when PIM state is disabled.

The standard SSM range is 232/8, but this can be extended to include any IP Multicast address with the Ethernet Routing Switch 55xx and 56xx implementation of SSM. Although the SSM range can be configured, configuring it for all multicast groups (224/4 or 224.0.0.0/240.0.0.0 or 224.0.0.0/255.0.0.0) is not allowed. The SSM range allows you to configure existing applications without changing their group configurations. This flexibility allows applications to take immediate advantage of SSM. Candidate RP and Static RP cannot be configured for the SSM group range. In SSM mode, group ranges outside the SSM range are processed as in PIM-SM mode i.e. the IGMP reports and PIM join/prune messages not in SSM range are processed just as in PIM-SM mode.

The system prohibits you from making a dynamic change in an SSM group range with existing multicast trees. You must disable PIM before you can make a change in the SSM group range. This procedure reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), it also causes an RP relearn delay of up to 60 seconds. This delay can be longer if the BSR is local.

Theory of operation

By default PIM-SSM is globally disabled. To start PIM-SSM, the user must configure the switch PIM mode to ssm. Configurations are persistently saved inside NVRAM.

The trigger for PIM-SSM operation is the receipt of (S,G) IGMP report. The DR on the receiver LAN then sends an (S,G) join towards the source. Each switch along the path to source determines the upstream from the route to source provided by unicast routing. Once the (S,G) join is received by the first-hop-router, data is forwarded downstream to all the subscribed receivers.

SSM only uses a subset of the PIM-SM features such as the shortest path tree, designated router (DR), and some messages (Hello, Join/Prune, and Assert). However, there are also some features that are unique to SSM. These features, which are described in the following sections, are extensions of the IGMP and PIM protocols.

PIM-SSM architecture requires routers to:

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

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

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

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

Static IP routing table

The static IP routing table provides the flexibility of separating the paths for unicast and multicast streams. This table is used only by the multicast routing protocols PIM-SM & PIM-SSM.

An entry in this static IP routing table has the following attributes:

- IP prefix / IP mask — denotes the destination network for which the route is being added.
- Reverse Path Forwarding (RPF) address — denotes the RPF neighbor towards the source.
- Route preference — the administrative distance for the given route.

*** Note:**

When the unicast routing table and the multicast-static IP routing table have different routes for the same destination network, then this administrative distance is compared with that of the protocol that contributed the route in unicast routing table. By providing administrative distance for every route, you have the flexibility to choose different distances for different networks.

- Route Status — can be enabled / disabled from CLI command.

Routes from the multicast-static ip routing table can not be redistributed. They are used only for RPF calculations in multicast protocols.

The following rules must be followed while determining reverse path forwarding:

- Direct / Local routes for a given destination take precedence over any route for the same destination in multicast-static IP routing table.
- If a route is present in the static table, and no route exists in the unicast routing table for the given destination, the route in the static table should be used.
- If a route is available in both the unicast routing table and also in the multicast static IP routing table, then the route from multicast static IP routing table is used only if its administrative distance is less than or equal to that of the unicast route entry.

*** Note:**

The comparison between unicast routing information and static mroute does not use prefix lengths.

- If no route exists in the multicast-static IP routing table for the given destination, then the route from the unicast routing table should be used, if available.
- Longest prefix match is performed when doing a lookup within the multicast-static IP routing table. The lookup ignores routes that are administratively disabled.

IGMP functionality over split multilink trunking (SMLT) network topologies

To support internet group management protocol over split multilink trunking, the following basic operations are performed:

- The groups on one SMLT aggregation switch are populated in the other SMLT aggregation switch over the IST to have a redundant path (Wiring Closet Switch). IGMP IST report messages are forwarded only to query ports and not IST links. If the report is received over SMLT, the switch then configures the SMLT links as the member ports for that group.

*** Note:**

If the IST group report message does not have SMLT ID information, then IST links are configured as the member ports for the group.

- IGMP queries received on an aggregation switch are flooded to all links (except the IST links) in the VLAN. IGMP queries reaching one aggregation switch are updated in the other aggregation switch using IST query messages. The aggregation switch uses the SMLT ID in the query message to configure the SMLT links as the querier ports. If an SMLT link goes down, the querier port moves to the IST link. This is to make sure that IGMP reports are always forwarded to the IGMP querier/PIM router.

*** Note:**

If the IGMP query is received on a link or MLT (not SMLT), the SMLT ID is set to zero in the message. The aggregation switch which receives the IST query message uses the SMLT ID in the query message to configure the SMLT links as the querier ports. If the “IST query message” does not have a SMLT ID (SMLT ID = 0), then the IST link is configured as the querier port.

- IGMP leave messages received on an aggregation switch are treated in the same way as IGMP reports. Leave messages are sent to the other aggregation switch over IST links in IST leave messages. The aggregation switch that receives the IST leave message removes the group membership for that member. If a switch has no more receivers for a group (receivers on SMLT links are exempted), switch sends an IST IGMP prune option. The switch that receives the IST IGMP leave with prune, removes the group membership of the IST link (if exists).
- IST messages are defined for different IGMP packet types. The different message types supported are:
 - IST IGMP Query message - contains the source IP address (querier IP address), VLAN, Group Address (in the case of group specific query), maximum response time and the SMLT ID (in the case IGMP queries are received on an SMLT).
 - IST IGMP Group Query message

- IST IGMP Group Report message - contains a source IP address, group address, VLAN, and SMLT.
- IST IGMP Group Leave message
- There are 2 types of SMLT messages that are processed by IGMP:
 - IGMP_PEER_SMLT_DOWN: The SMLT peer that receives this message sends IST IGMP group reports to the other peer for all members of that specific SMLT. The SMLT id is set to 0, so all reports are learned on the other peer's IST port (also applicable to the querier if it is located on a port member).
 - IGMP_PEER_SMLT_UP: Same as IGMP_PEER_SMLT_DOWN, however the SMLT id is set to the proper value.
- If an SMLT peer receives an SMLT IGMP message and the SMLT id specified in that message is down, it programs the IST port instead. The IST has to be up and that particular SMLT has to be configured.

IGMP Multicast flood control

IGMP Multicast flood control limits IP multicast traffic without inhibiting other control protocols. By minimizing IP multicast flooding in the network, it eliminates the necessity of queries sent by the switch when IGMP snooping is enabled.

When enabled, IGMP multicast flood control detects and limits sending multicast streams to multicast router ports (static & dynamic) when no clients are detected by redirecting native multicast streams to the CPU via an installed hardware filter.

 **Note:**

This is a global feature and applied to all VLANs. By default, this feature is disabled.

 **Important:**

The IGMP Multicast flood control feature is not available on ERS 5510 devices.

Chapter 5: IPv6 routing fundamentals

This chapter provides an introduction to IPv6 routing and the IPv6 routing features supported on the Avaya Ethernet Routing Switch 5000 Series.

IPv6 routing features

The Ethernet Routing Switch 5000 Series provides support for configurable IPv6 static routes and per-VLAN IPv6 routes.

Supported features include the following:

- Multiple configurable IPv6 interfaces associated with VLANs
- One IPv6 global addresses (automatically inserted into the routing table) per IPv6 interface
- Multiple configurable static route entries in the IPv6 routing table
- Router functionality based on the routing table constructed using the two preceding methods listed

The ERS 5510, 5520 and 5530 do not support IPv6 routing, non-local IPv6 static routes, nor do hybrid stacks containing those units.

The scaling limits on the routing table are as follows.

- 512 maximum IPv6 static routes
- 256 maximum IPv6 interfaces

Host autoconfiguration

The Avaya Ethernet Routing Switch 5000 Series supports the Neighbor Discovery Protocol for IPv6. Using Router Advertisements forwarded by the switch, hosts can perform stateless autoconfiguration of site-local and global IPv6 addresses.

Stateless autoconfiguration enables serverless basic configuration of IPv6 hosts.

With stateless autoconfiguration, the IPv6 address is created as follows:

autoconfigured IPv6 address = network prefix + IPv6 Interface identifier

To create the IPv6 Interface identifier, stateless autoconfiguration uses a modified Extended Unique Identifier (EUI-64) format derived from the interface MAC address.

The modified EUI-64 information is created from the 48-bit (6-byte) MAC address as follows:

1. Hexadecimal digits 0xff-fe are inserted between the third and fourth bytes of the MAC address to obtain an EUI-64 address.
2. The universal/local bit, the second lower-order bit of the first byte of the EUI-64 address, is complemented (changed from zero to one).

For example, host A uses the MAC address 00-AA-00-3F-2A-1C. The following steps show how this MAC address can be converted to modified EUI-64 format for use in an IPv6 address:

1. Given the MAC address:

00-AA-00-3F-2A-1C

Convert it to an EUI-64 address by inserting 0xFFFE between the third and fourth bytes:

00-AA-00-FF-FE-3F-2A-1C

2. Complement the Universal/Local (U/L) bit.

The first byte in binary form is 00000000. When the seventh bit (universal/local bit) is complemented, it becomes 00000010 (0x02).

In this case, the result is:

02-AA-00-FF-FE-3F-2A-1C

OR

2AA:FF:FE3F:2A1C

Upon initialization, hosts use the common link-local prefix FE80 to autoconfigure a link-local address.

In this example, the link-local address for host A with the MAC address 00-AA-00-3F-2A-1C is FE80::2AA:FF:FE3F:2A1C. Because the common FE80 prefix is used for link-local addresses, a router is not required for link-local address autoconfiguration. Before using the autoconfigured link-local address, the host performs a check using the Neighbor Discovery Protocol to ensure that its autoconfigured address is not a duplicate address.

For autoconfiguration of site-local and global addresses, a router must be present in the network. In these cases, stateless autoconfiguration uses the following format:

autoconfigured IPv6 address = network prefix (from router advertisement) + IPv6 Interface identifier

To create the IPv6 address, stateless autoconfiguration uses the network prefix information in the router advertisement messages. The modified Extended Unique Identifier (EUI-64) format provides the remaining address.

For example, host A with MAC address 00-AA-00-3F-2A-1C, combined with network prefix 2001::/64 provided by router advertisement, uses an IPv6 address 2001::2AA:FF:FE3F:2A1C.

You can use the `ipv6 nd prefix` command to specify the prefixes to advertise in the router advertisement messages.

The following are the states associated with autoconfiguration addresses:

- **Tentative:** the address is being verified as unique (link-local address)
- **Valid:** an address from which unicast traffic can be sent and received and can be in one of two states
- **Deprecated:** an address that remains valid but is withheld for new communication
- **Preferred:** an address for which uniqueness was verified for unrestricted use
- **Invalid:** an address for which a node can no longer send or receive unicast traffic

A valid lifetime is the length of time of the preferred and deprecated state. The preferred lifetime is the length of time for the tentative, preferred, and deprecated state.

IPv6 static routes

Static routes provide a method for establishing route reachability. This function provides routing information from the forwarding database to the forwarding plane. Only enabled static routes are submitted to the Route Table Manager (RTM), which determines the best route based on reachability, route preference, and cost. The RTM communicates all updates to best routes to the forwarding plane.

Avaya does not recommend IPv6 static routes over an SMLT/IST setup.

You can configure the following options when configuring a static route:

- **next hop:** Specifies the next hop to the destination address. Configure a static route either with a next hop that exists on a locally attached network or a next hop that is reachable through a default static route. The static route is available as long as the next hop is reachable.
- **cost:** Specifies the cost or distance ratio to reach the destination address. The switch prefers lower-cost routes over higher-cost routes.
- **Route preference:** Specifies the preference value associated with a particular route. The switch prefers routes with lower preferences over those with higher preferences. Whereas the cost value assigns an administrative weight to the route itself, the preference generally

assigns a weight to the process used to discover the route (for example, by static route rather than by dynamic protocol such as RIP).

- administrative status: controls when the static route is considered for forwarding. Administrative status differs from the operational status. An admin-enabled static route can still be unreachable and not be used for forwarding. An admin-disabled static route is operationally a nonexistent route.

To configure a default static route, enter a value of 0::0 for the prefix and 0 for the prefix length.

Events that affect static route operation include user-configured changes or other system events. The table below describes these changes.

Table 8: Static route operation changes

Action	Result
Disabling the administrative status of the static route	Makes the static route unavailable for forwarding.
Deleting the IPv6 addresses of a VLAN	Permanently deletes the static routes with the corresponding local neighbors from the RTM, the forwarding database, and the configuration database.
Deleting a VLAN	Removes static routes with a local next-hop option from the configuration database. Static routes with a nonlocal next-hop option become inactive (they are removed from the forwarding database).
Disabling forwarding on a VLAN	Static routes reachable through the locally attached network become inactive.
Disabling a VLAN	Makes the static routes inactive.
Disabling IPv6 forwarding globally	Stops the forwarding of all IPv6 traffic.
Learning changes about a dynamically learned neighbor	When a neighbor becomes unreachable or is deleted, the static route with the neighbor becomes inactive, and the configuration is not affected. When the neighbor becomes reachable, the static route with the neighbor becomes active in the configuration and is added to the RTM and forwarding database.
Enabling a static route	Adds the route to the RTM to change certain static routes to active.
Deleting a static route	Permanently deletes a static route from the configuration.
Disabling a static route	Stops traffic on the static route but does not remove the route from the configuration.

Action	Result
Deleting or disabling a tunnel	Deletes or disables a tunnel and removes the tunnel entry from the forwarding table.
Enabling the tunnel	Enables a tunnel, activates the tunnel static routes and adds an entry to the forwarding table.

To provide stability and load balancing, you can specify alternative paths to the same destination with multiple static routes. You can enter multiple routes (for example, multiple default routes) that use different costs and the lowest cost route that is reachable is the one that appears in the routing table. If you enter multiple next hops for the same route with the same cost, the switch does not replace the existing route.

If you enter the same route with the same cost and a different next hop, the first route is used. However, if that first route becomes unreachable, the second route (with a different next hop) is activated with no connectivity loss.

Management route

The static route configured for the management VLAN is maintained separately from the routing table. Because traffic on the switch management VLAN refers to this route first before checking the routing table, the switch management traffic can be incorrectly forwarded from the management VLAN, even though a specific route exists in the routing table. Configure your management route carefully to take into account this potential issue.

IPv6 DHCP Relay

IPv6 DHCP Relay for the Ethernet Routing Switch 5000 Series allows the routing switch to act as an IPv6 DHCP (or DHCPv6) relay agent, as described in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*.

A DHCPv6 relay agent is used to relay messages between a DHCPv6 client and a DHCPv6 server connected to different VLANs.

DHCP for IPv6 enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. DHCP supports automatic allocation of reusable network addresses and of additional configuration parameters.

In basic DHCP operation, a client locates and communicates with a DHCP server using a reserved, link-scoped multicast address. For this to be possible, the client and the server have to be connected to the same link.

To request the assignment of one or more IPv6 addresses, a client first locates a DHCP server and then requests the assignment of addresses and other configuration information from the server. The client sends a Solicit message to the All_DHCP_Relay_Agents_and_Servers (FF02::1:2) multicast address to find available DHCP servers. Any server that can meet the requirements from the client responds with an Advertise message. The client then chooses one of the servers and sends a Request message to the server asking for confirmed assignment of addresses and other configuration information. The server responds with a Reply message that contains the confirmed addresses and configuration.

IPv6 DHCP clients use link-local addresses to send and receive DHCP messages.

However, in some situations, for ease of management, economy or scalability, it can be desirable to allow a DHCP client to communicate with a DHCP server that is not connected to the same link. The DHCP relay agent makes this possible, relaying the messages between the client and the remote server.

To allow a DHCP client to send a message to a DHCP server that is not attached to the same link, you must configure a DHCP relay agent on the client link to relay messages between the client and server. The operation of the relay agent is transparent to the client.

A relay agent relays messages from clients and from other relay agents.

IPv6-in-IPv4 tunnels

The IPv6 in IPv4 tunneling feature enables isolated IPv6 sites to communicate with other IPv6 sites by encapsulating IPv6 packets in IPv4 packets through an IPv4 network.

The following figure shows an example of an IPv6-in-IPv4 tunnel.

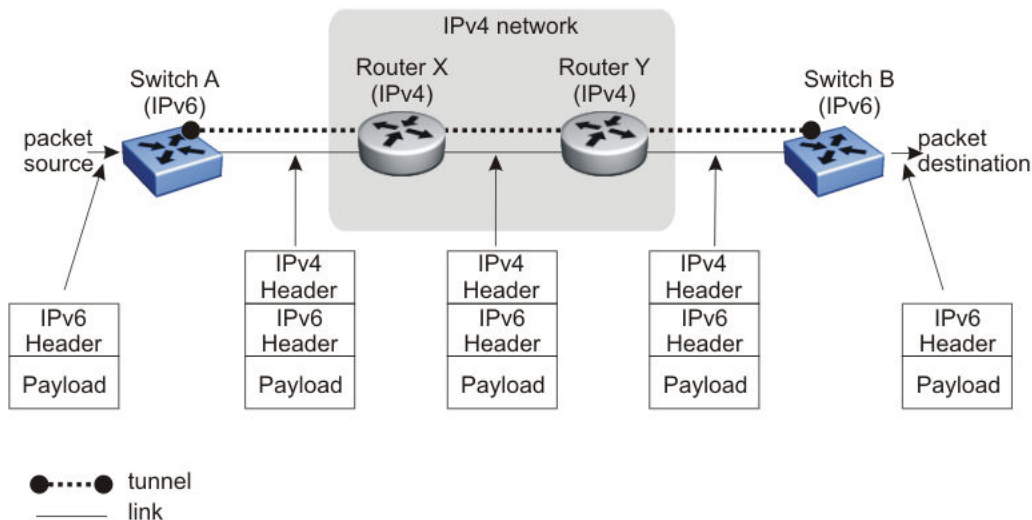


Figure 32: IPv6-in-IPv4 tunnel example

In the preceding figure, Switch A is the entry node of the tunnel (encapsulating node), and Switch B is the exit node of the tunnel (decapsulating node).

1. Switch A receives the IPv6 packet from the source and determines that it must be forwarded out the tunnel interface.
2. Switch A encapsulates the IPv6 packet in an IPv4 header and transmits the encapsulated packet.

The source address in the IPv4 header is the IPv4 address of the local tunnel interface on switch A. The destination address is the IPv4 address of the remote tunnel interface on switch B.
3. Using the IPv4 header, the intermediate IPv4 routers forward the encapsulated packet through the IPv4 network to switch B.
4. Switch B receives the IPv4 packet, removes the outer IPv4 header, and then processes the decapsulated IPv6 packet.

The Ethernet Routing Switch 5000 Series supports manually configured tunnels. To enable the tunnel, you must manually specify the IPv4 addresses of the local and remote endpoints of the tunnel.

Related topics:

[Tunneling limitations](#) on page 128

Tunneling limitations

The following limitations apply to IPv6-in-IPv4 tunnels:

- Routing of data packets on these configured tunnels is not supported; that is, only applications originating or terminating on the Ethernet Routing Switch 5000 Series can use these tunnels.
- The maximum number of supported tunnels is four.
- Dynamic IPv6 routing protocols are not supported on tunnel interfaces.
- IPv4 path MTU discovery is not supported for data on tunnels.
- IPv4 fragmentation and reassembly is not supported over tunnels.
- IPv4 ICMP errors are not translated to IPv6 ICMP errors.
- IPv6 tunneling is not supported in Layer 2 mode.
- When the management VLAN IP is configured as a tunnel end point, tunnel functionality is not supported after a stack transition (from stack to switch or switch to stack) because the IPv4 source address in the stack and switch are different. In this case, you must reconfigure the tunnel source to the new management VLAN IP address.
- All ACLI commands for IPv6 can only be executed on the base unit of a stack.
- Stacking support is provided only on pure 5500 or 5600 stacks or on Hybrid stacks.

IPv6 management

This module provides information about the IPv6 management feature of the Avaya Ethernet Routing Switch 5000 Series switch platforms.

Chapter 6: IP routing configuration using ACLI

This chapter describes the procedures you can use to configure routable VLANs using Avaya Command Line Interface (ACLI).

The Avaya Ethernet Routing Switch 5000 Series are Layer 3 (L3) switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address and MAC address are 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.

For more information on creating and configuring VLANs, see *Configuration — VLANs, Spanning Tree, and Link Aggregation* (NN47200-502).

IP routing configuration procedures

To configure inter-VLAN routing on the switch, perform the following steps:

1. Enable IP routing globally.
2. Assign an IP address to a specific VLAN or router port.

Routing is automatically enabled on the VLAN or router port when you assign an IP address to either of these interfaces.

In the preceding procedure, you are not required to enable IP routing as the first step. You can configure all IP routing parameters on the Avaya Ethernet Routing Switch 5000 Series before you enable routing for the switch.

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

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip routing
```

Variable definitions

The following table describes the 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 blocking on the switch.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip routing
```

Configuring an IP address for a VLAN

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

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Variable definitions

The following table describes the 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

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Variable definitions

The following table describes the command variables.

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

Configuring a secondary IP address for a VLAN

Use this procedure to configure a secondary IP interface to a VLAN (also known as multinetting). You can have a maximum of eight secondary IP addresses for every primary address, and you must configure the primary address before configuring any secondary addresses.

Primary and secondary interfaces must reside on different subnets.

To remove a primary IP address from a VLAN, you must first remove all secondary addresses from the VLAN.

Prerequisites

- Configure a primary IP address on the VLAN.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip address <ip address> <mask> [<mac offset>] secondary
```

Variable definitions

The following table describes the command variables.

Variable	Value
no	Removes the configured IP address. To remove a primary IP address from a VLAN, you must first remove all secondary addresses from the VLAN.
<ipaddr>	Specifies the IP address to attach to the VLAN.
<mask>	Specifies the subnet mask to attach to the VLAN

Variable	Value
[<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.

Job aid: Example of adding a secondary IP interface to a VLAN

Primary and secondary interfaces must reside on different subnets. In the following example, 4.1.0.10 is the primary IP and 4.1.1.10 is the secondary IP.

```
(config)# interface vlan 4
(config-if)# ip address 4.1.0.10 255.255.255.0 6
(config-if)# ip address 4.1.1.10 255.255.255.0 7 secondary
```

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.

Procedure steps

1. Log on to the Privileged EXEC Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
show vlan ip [id <vid>] [vrf {vrfName}] [vrfids {vrf_ids}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[id <vid>]	Specifies a list or range of VLAN IDs to be displayed. VLAN IDs range from 1 to 4094..
[vrf {vrfName}]	Specifies the VRF instance identified by the alphanumeric identifier (name) for which to display VLAN information.

Variable	Value
[vrfids {vrf_ids}]	Specifies the VRF instance(s) identified by the numerical ID(s) for which to display route information.

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.
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 in the routing table.

Route entries appear in ascending order of the destination IP addresses.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip route [<dest-ip>] [-s <subnet> <mask>] [summary] [vrf
{vrfName}] [vrfids {vrf_ids}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[<dest-ip>]	Specifies the destination IP address of the route to display.
[-s <subnet> <mask>]	Specifies the destination subnet of the routes to display.
[summary]	Displays a summary of IP route information.
[vrf {vrfName}]	Specifies the VRF instance identified by the alphanumeric identifier (name) for which to display route information.
[vrfids {vrf_ids}]	Specifies the VRF instance(s) identified by the numerical ID(s) for which to display route information.

Job aid

The following show sample outputs for the **show ip route** command.

show ip route command output

Ip Route								
DST	MASK	NEXT	COST	VLAN	PORT	PROT	TYPE	PRF
0.0.0.0	0.0.0.0	10.3.2.13 7	1	1	1/21	S	IB	5
2.2.2.0	255.255.255 .0	2.2.2.2	1	2	----	C	DB	0
10.3.2.0	255.255.255 .0	10.3.2.19 9	1	1	----	C	DB	0
Total Routes: 3								
TYPE Legend: I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route, U=Unresolved Route, N=Not in HW								

show ip route summary command output

```

-----
Connected routes :      65
Static routes    :       2
RIP routes       :     512
OSPF routes      :     512
BGP routes       :       9
-----
Total routes     :    1100
-----

```

Performing a traceroute

Use this procedure to display the route taken by IP packets to a specified host.

When applied to a stack, this procedure can be executed only on the base unit.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```

traceroute <Hostname|A.B.C.D.|ipv6-addr> <-m> <-p> <-q> <-v>
<-w> <1-1464>

```

*** Note:**

You can enter CTRL+C at any time to interrupt the command.

Variable definitions

The following table describes the command variables.

Variable	Value
Hostname	Specifies the name of the remote host.
A.B.C.D.	Specifies the IP address of the remote host.
ipv6-addr	Specifies the IPv6 address of the remote host.
-m	Specifies the maximum time to live (ttl). The value for this parameter is in the range from 1-255. The default value is 10. Example: traceroute 10.3.2.134 -m 10

Variable	Value
-p	Specifies the base UDP port number. The value for this parameter is in the range from 0-65535. Example: traceroute 1.2.3.4 -p 87
-q	Specifies the number of probes per time to live. The value for this parameter is in the range from 1-255. The default value is 3. Example: traceroute 10.3.2.134 -q 3
-v	Specifies verbose mode. Example: traceroute 10.3.2.134 -v
-w	Specifies the wait time per probe. The value for this parameter is in the range from 1-255. The default value is 5 seconds. Example: traceroute 10.3.2.134 -w 15
<1-1464>	Specifies the UDP probe packet size. TIP: probe packet size is 40 plus specified data length in bytes. Example: traceroute 10.3.2.134 -w 60

Entering Router Configuration mode

Use this procedure to enter Router Configuration mode and configure parameters related to routing protocols. Router Configuration mode is used to configure RIP, OSPF, VRRP, BGP and VRF.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
router {rip | ospf | vrrp | bgp | vrf {vrfName}}
```

Accessing Loopback Interface Configuration mode

Access Loopback Interface Configuration mode to configure a circuitless IP (CLIP) interface.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
interface loopback {1-16}
```

Variable definitions

The following table describes the parameters for the `interface loopback` command.

Variable	Value
{1–16}	Specifies the loopback interface identifier. Values range from 1 to 16.

Configuring a CLIP interface

Configure a circuitless IP (CLIP) interface to provide a virtual interface that is not associated with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to the switch.

! Important:

You can configure a maximum of 16 CLIP interfaces on each ERS 5000 Series device.

Prerequisites

- Enable IP routing globally.

Procedure steps

1. Log on to the Loopback Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip [address <A.B.C.D> / <mask>] [area <A.B.C.D>] [ospf]
```

Variable definitions

The following table describes the parameters for the `ip` command.

Variable	Value
address <A.B.C.D> / <mask>	Specifies the CLIP interface IP address and subnet mask.
area <A.B.C.D>	Assigns the CLIP interface to a specific area.
ospf	Enables OSPF on the CLIP.

Variable	Value
	<p>! Important:</p> <p>OSPF runs only in passive mode on a CLIP interface.</p>

Deleting CLIP configuration parameters

Use this procedure to clear or delete CLIP configuration parameters from a loopback interface.

Procedure steps

1. Log on to the Loopback Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
no ip [address <A.B.C.D> / <mask>] [area] [ospf]
```

Variable definitions

The following table describes the parameters for the `no ip` command.

Variable	Value
address <A.B.C.D> / <mask>	Deletes the CLIP IP address and subnet mask.
area	Removes the CLIP from a specific area.
ospf	Disables OSPF on the CLIP.

Restoring CLIP to default

Use this procedure to restore CLIP configuration parameters for a loopback interface to default values.

Procedure steps

1. Log on to the Loopback Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
default ip [area] [ospf]
```

Variable definitions

The following table describes the parameters for the `default ip` command.

Variable	Value
area	Removes the CLIP from a specific area.
ospf	Disables OSPF on the CLIP.

Displaying CLIP information

Use this procedure to display and verify CLIP configuration information for a switch.

Procedure steps

1. Log on to the Privileged EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show interface loopback [1-16]
```

Variable definitions

The following table describes the parameters for the `show interface loopback` command.

Variable	Value
1-16	<p>Displays CLIP information for a specific loopback interface. Values range from 1 to 16.</p> <p>* Note:</p> <p>If you do not include this variable, the switch displays information for all configured CLIPs.</p>

Chapter 7: IP routing configuration using EDM

This chapter describes the procedures you can use to configure routable VLANs using Enterprise Device Manager (EDM).

The Avaya Ethernet Routing Switch 5000 Series are Layer 3 (L3) switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address and MAC address are attached to the VLAN. When you enable routing 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 IP routing using EDM

Use the following procedure to configure IP routing on VLANs.

1. Enable IP routing globally.
2. Assign an IP address to a specific VLAN.

For the preceding procedure, you are not required to enable IP routing as the first step. You can configure all IP routing parameters on the Avaya Ethernet Routing Switch 5000 Series before you enable routing for globally the switch.

Configuring global IP routing status and ARP lifetime

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

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Globals** tab.

4. Select **forwarding** in the Forwarding field to enable routing.
5. Enter the ARP life time in the **ARPLifeTime** field.
6. On the toolbar, click **Apply**.
7. On the toolbar, you can click **Refresh** to verify the global IP routing and ARP lifetime configuration.

Variable definitions

The following table describes the fields of the **Globals** tab.

Field	Description
Forwarding	Indicates whether routing is enabled (forwarding) or disabled (nonforwarding) 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. Range is 1–255. Default value is 64 seconds. This is a read only field.
ReasmTimeout	Indicates the maximum number of seconds that received fragments are held while they await reassembly at this entity. Default value is 60 seconds. This is a read only field.
ARPLifeTime	Specifies the lifetime in minutes of an ARP entry within the system. Range is 5-360. Default is 360 minutes.
AdminEnabled	Enables and disables forwarding next hop.
OperEnabled	A read only field indicating the current operational status of forwarding next hop: true (enabled) or false (disabled).
DirectedBroadcast	Enables and disables IP directed broadcast

Configuring IP directed broadcasts per VLAN

Use this procedure to configure IP directed broadcasts on a VLAN basis.

Procedure steps

1. From the navigation pane, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.

3. In the **Basic** tab, select the VLAN ID that you want to configure with directed broadcast.
4. On the toolbar, click **IP**.
5. In the work area, click **DirectedBroadcast**.
6. Select the **DirectedBroadcast** checkbox to enable, or clear the checkbox to disable.
7. On the toolbar, click **Apply**.

Configuring an IP address and enabling routing for a VLAN

Use the following procedure to configure an IP address and enable routing for a VLAN.

Prerequisites

- Enable routing globally on the switch.


Procedure steps

1. From the navigation pane, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the work area, select the desired VLAN from the table.
4. In the toolbar, click **IP**.
5. In the toolbar, click **Insert**.
6. Configure the parameters as required.
7. Click **Insert**.

To enable or disable IP routing for a specific VLAN, go to the VLANs page and select the Routing option. A value of true enables IP routing for the VLAN, a value of false disables IP routing for the VLAN.

Variable definitions

The following table describes the fields of the **IP Address** tab.

Field	Description
IpAddress	Specifies the IP address to associate with the selected VLAN.
NetMask	Specifies the subnet mask.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface.
ReasmMaxSize	Specifies the size of the largest IP datagram which this entity can reassemble from fragmented incoming IP datagrams received on this interface.
VlanId	Specifies the VLAN ID.
MacOffset	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.
SecondaryIf	<p>Indicates whether or not this entry corresponds to a secondary interface. If the value is false, then this is the primary IP address, if the value is true, then this is a secondary IP address.</p> <p> Note: You can assign 1 primary IP address and up to 8 secondary IP addresses to a VLAN.</p>

Displaying configured IP Addresses

Use the following procedure to display configured IP addresses on the switch.

Procedure steps

1. From the navigation pane, double-click **IP Routing**.
2. In the IP Routing tree, click **IP**.
3. In the work area, click the **Addresses** tab to view configured IP addresses on the switch.

Variable definitions

The following table describes the fields of the **Addresses** tab.

Field	Description
IfIndex	Specifies the name of the VLAN.
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 number. 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.
SecondaryIf	Indicates whether or not this entry corresponds to a secondary interface. If the value is false , then this is the primary IP address, if the value is true , then this is a secondary IP address.

Configuring a CLIP interface

Configure a circuitless IP (CLIP) interface to provide a virtual interface that is not associated with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to the switch.

Important:

You can configure a maximum of 16 CLIP interfaces on each ERS 5000 Series device.

Prerequisites

- Enable IP routing globally.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Circuitless IP** tab.
4. On the toolbar, click **Insert**.
5. Configure the CLIP interface as required.
6. Click **Insert**.
7. On the toolbar, click **Refresh** to verify the CLIP interface configuration.

Variable definitions

Variable	Value
IfIndex	Specifies the identifier of loopback interface on which to configure CLIP. Values range from 1 to 16.
IpAddress	Specifies the CLIP IP address.
NetMask	Specifies the CLIP IP subnet mask.

Deleting a CLIP interface

Use this procedure to delete CLIP from a loopback interface.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Circuitless IP** tab.
4. In the Circuitless IP work area, click the **IfIndex** of the CLIP to delete.
5. On the toolbar, click **Delete**.
6. On the toolbar, click **Refresh** to verify the CLIP interface is deleted from the system.

Configuring a CLIP interface for OSPF

Use this procedure to configure a CLIP interface to run OSPF.

! **Important:**

OSPF runs only in passive mode on a CLIP interface.

Prerequisites

- Enable IP routing globally.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Circuitless IP** tab.
4. In the Circuitless IP work area, click the **IfIndex** of a CLIP.
5. On the toolbar, click **OSPF**.
6. Configure OSPF for the CLIP interface.
7. On the toolbar, click **Apply**.
8. On the toolbar, click **Refresh** to verify the OSPF configuration for the CLIP interface.

Variable definitions

Variable	Value
Enable	Enables (selected) or disables (cleared) OSPF for the CLIP interface.
IfAreald	Assigns the CLIP to a specific area.

Chapter 8: VRF Lite configuration using ACLI

This chapter provides information for creating and managing Virtual Router Forwarding (VRF) Lite, using the Avaya Command Line Interface (ACLI).

! Important:

VRF Lite is supported on the ERS 5600 series only.

Configuring a VRF instance

Use this procedure to create and configure a new VRF instance or modify an existing VRF instance.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip vrf {vrfName} [max-routes {0-4000}] [vrfid {1-3}] [name <name>]
```

Variable definitions

The following table describes the parameters for the `ip vrf` command.

Variable	Value
{vrfName}	Specifies the alphanumeric identifier (name) assigned to the VRF instance. Value ranges from 1 to 16 characters.
max-routes{0-4000}	Specifies a maximum number of routes for the VRF instance. Values range from 0 to 4000 for the default VRF and 0 to 500 for non-default VRFs. Note: a maximum of 4000 routes exists between all the VRF instances. A default of 500 max-routes is available when non-

Variable	Value
	default VRFs are created. If max-routes=0 , no L3 interfaces can be created in that VRF instance.
vrfid{1–3}	Specifies a numerical ID for the VRF instance. Values range from 1 to 3.
name<name>	Changes the alphanumeric identifier for the VRF instance to a new value. Value ranges from 1 to 16 characters.

Deleting a VRF instance

Use this procedure to delete an existing VRF instance from your system.

! Important:

Deleting a VRF instance, will remove all L3 interfaces and its associated L3 configurations that is, static arps, static routes from that VRF instance. The VLAN associated with the VRF will be moved to default VRF.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
no ip vrf <vrfName>
```

Variable definitions

The following table describes the parameters for the `no ip vrf` command.

Variable	Value
<vrfName>	Specifies the alphanumeric identifier of the VRF instance to delete.

Displaying VRF instances

Use this procedure to display and verify the configuration for VRF instances on the switch.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:


```
show ip vrf [max-routes] [vrfid {1-3}] [vrfName]
```

Variable definitions

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

Variable	Value
max-routes	Displays the maximum number of routes configured for the VRF instance. Values range from 0 to 4000.
vrfid{1-3}	Displays configuration information for a specific VRF instance numerical ID.
vrfName	Displays configuration information for a specific VRF instance name.
<p>* Note:</p> <p>If you do not include variables with the <code>show ip vrf</code> command, the switch displays complete configuration information for all VRF instances.</p>	

Associating a VLAN with a VRF instance

Use this procedure to associate a VLAN with a specific VRF instance.

! Important:

You cannot associate a VLAN with a VRF instance if the VLAN has an IP address configured. You can configure the VLAN IP address after you associate the VLAN with the VRF instance.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:


```
vrf {vrfName}
```

Variable definitions

The following table describes the parameters for the `vrf` command.

Variable	Value
<code>{vrfName}</code>	Specifies the alphanumeric identifier (name) assigned to the VRF instance.

Removing a VLAN association with a VRF instance

Use this procedure to remove the association of a VLAN with a specific VRF instance.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
no vrf
```

Displaying VLAN and VRF instance associations

Use this procedure to display information about VLAN and VRF instance associations within your system.

Procedure steps

1. Log on to the Privileged EXEC mode in ACLI.
2. To display a list of configured VLANs and associated VRF instances, enter the following command:
3. To display detailed information for VLANs associated with a specific VRF instance name, enter the following command:

```
show vlan vrf
```

```
show vlan vrf {vrfName}
```

4. To display detailed information for VLANs associated with a specific VRF instance ID, enter the following command:

```
show vlan vrfids {0-3}
```

Variable definitions

The following table describes the variables available with the `show vlan vrf` and `show vlan vrfids` commands.

Variable	Value
<code>{vrfName}</code>	Displays detailed configuration information for a VLAN associated with a specific VRF instance, identified by name.
<code>{0-3}</code>	Displays detailed configuration information for a VLAN associated with a specific VRF instance, identified by ID.

Associating a port with a VRF instance

Use this procedure to associate a port with a specific VRF instance.

! Important:

To configure a Brouter port on a non-default VRF, associate the port to a VRF, then configure the port as a Brouter.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
vrf {vrfName}
```

Variable definitions

The following table describes the parameters for the `vrf` command.

Variable	Value
<code>{vrfName}</code>	Specifies the alphanumeric identifier (name) assigned to the VRF instance.

Removing a port association with a VRF instance

Use this procedure to remove the association of a port with a specific VRF instance.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
no vrf
```

Displaying switch port and VRF instance associations

Use this procedure to display information about switch port and VRF instance associations within your system.

Procedure steps

1. Log on to the Privileged EXEC mode in ACLI.
2. To display a list of configured switch ports and associated VRF instances, enter the following command:

```
show int vrf <portlist>
```

Variable definitions

The following table describes the parameters for the `show int vrf` command.

Variable	Value
<code><portlist></code>	Displays VRF instance association information for a specific switch port or list of ports.

Changing the router mode to a VRF instance

Use this procedure to change the router configuration mode for a particular non-default VRF instance.

All commands executed under this VRF mode will be applied to that particular VRF.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
router vrf {vrfName}
```

Variable definitions

The following table describes the parameters for the `router vrf` command.

Variable	Value
{vrfName}	Specifies the alphanumeric identifier (name) assigned to the non-default VRF instance.

Chapter 9: VRF Lite configuration using EDM

This chapter provides information for creating and managing Virtual Router Forwarding (VRF) Lite, using Enterprise Device Manager (EDM).

! Important:

VRF Lite is supported on the ERS 5600 series only.

Configuring a VRF instance

Use this procedure to modify the configuration parameters for an existing VRF instance.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **VRF**.
3. In the VRF work area, double-click an **Id** cell, to modify the configuration for that parameter.
4. On the toolbar, click **Apply**.
5. On the toolbar, click **Refresh** to verify the VRF instance configuration.

Variable definitions

Variable	Value
Id	Specifies a numerical ID for the VRF instance. Values range from 1 to 3.
Name	Specifies the alphanumeric identifier (name) assigned to the VRF instance.
MaxRoutes	Specifies a maximum number of routes for the VRF instance. Values range from 0 to 500. Note: a maximum of 4000 routes exists between all the VRF instances. A default of

Variable	Value
	500 max-routes is available when non-default VRFs are created. If max-routes=0 , no L3 interfaces can be created in that VRF instance.

Creating a VRF instance

Use this procedure to create a new VRF instance.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **VRF**.
3. On the VRF work area toolbar, click **Insert**.
4. Configure the new VRF instance as required.
5. Click **Insert**.
6. On the toolbar, click **Refresh** to verify the VRF instance configuration.

Variable definitions

Variable	Value
Id	Specifies a numerical ID for the VRF instance. Values range from 1 to 3.
Name	Specifies the alphanumeric identifier (name) to assign to the VRF instance.
MaxRoutes	Specifies a maximum number of routes for the VRF instance. Values range from 0 to 500.

Deleting a VRF instance

Use this procedure to delete an existing VRF instance.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **VRF**.
3. In the VRF work area, click the **Id** of the VRF instance to delete.
4. On the toolbar, click **Delete**.
5. On the toolbar, click **Refresh** to verify the VRF instance configuration.

Viewing brouter port and VRF associations

Use this procedure to view each port and associated VRFs. You can also change the VRFs associated with the port if the port has no IP address.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **VRF**.
3. In the VRF work area, click the **VRF-Ports** tab.
4. To change the VRF, double-click the **BrouterVrflid** field for the port.

+ Tip:

You can associate a port with more than one VRF.

5. Choose the required VRFs, and click **Ok**.
6. Click **Apply**.

Variable definitions

Variable	Value
rcPortIndex	Specifies the port number.
Type	Specifies the port type.
BrouterVrflid	Specifies the VRF ID for this brouter port.
BrouterVrfname	Specifies the VRF name for this brouter port.

Associating a VLAN with a VRF instance

Use this procedure to associate a VLAN with a specific VRF instance.

! Important:

You cannot associate a VLAN with a VRF instance if the VLAN has an IP address configured. You can configure the VLAN IP address after you associate the VLAN with the VRF instance.

Procedure steps

1. From the navigation pane, double-click **VLAN**.
2. In the VLAN tree, click **VLANS**.
3. In the VLANs work area, click the **Basic** tab.
4. Double-click the **Vrflid** cell for the VLAN to assign a VRF instance.
5. Type a value.
6. On the toolbar, click **Apply**.
7. On the toolbar, click **Refresh** to verify the VLAN is associated with the VRF instance.

Variable definitions

Variable	Value
Vrflid	Indicates the numerical identifier assigned to the VRF instance. Value ranges from 0 to 3.

Associating a switch port with a VRF instance

Use this procedure to associate a switch port with a specific VRF instance.

! Important:

You cannot associate a switch port with a VRF instance if the port has an IP address configured. You can configure the port IP address after you associate the port with the VRF instance.

Procedure steps

1. In the Device Physical View, select a switch port.
2. From the navigation pane, double-click **Edit**.
3. In the Edit tree, click **Chassis**.
4. In the Chassis tree, click **Ports**.
5. In the Ports work area, click the **VRF** tab.
6. Click the ellipsis (...) to the right of the **BrouterVrfid** box.
7. Select a VRF instance.
8. Click **Ok**.
9. On the toolbar, click **Apply**.
10. On the toolbar, you can click **Refresh** to verify the port association with the VRF instance.

Variable definitions

Variable	Value
BrouterVrfid	Indicates the numerical ID of the VRF instance currently associated with the port. Values range from 1 to 3. You can change the VRF instance associated with the port by clicking the ellipsis. A value of 0 indicates that the port is associated with default VRF.
BrouterVrfName	Indicates the name of the VRF instance associated with the port.

Selecting and launching a VRF context view

Use this procedure to switch to another VRF context view when you use EDM. GlobalRouter is the default view at log in. You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view. You can open only five tabs for each EDM session.

! Important:

If you log out from the GRT view, the system generates a warning: all tabs close and your session terminates. If you close a VRF view tab, you close only that view.

*** Note:**

The Set VRF Context view function is not available to users in a service provider deployment where only a tenant VRF view is assigned. If you use a tenant VRF view, Avaya recommends that you use the applicable EDM plugin with COM to access EDM. COM provides VRF mapping and Role-Based Access Control.

Procedure steps

1. From the navigation pane, double-click **VRF Context View**.
2. In the VRF Context View tree, click **Set VRF Context View**.
3. In the Set VRF Context View work area, click the **VRF** tab.
4. Select a context to view.
5. Click **Launch VRF Context view**.

Variable definitions

Variable	Value
Id	Displays the unique VRF ID.
Name	Displays the name of the virtual router.

Chapter 10: Brouter port configuration using ACLI

A brouter port is a single-port VLAN that can route IP packets and bridge all non-routable traffic. This chapter describes procedures you can use to configure a brouter port using Avaya Command Line Interface (ACLI).

Configuring a brouter port

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

Procedure steps

! Important:

To configure brouter port for a non-default VRF a port should be first associated to a VRF and then brouter should be configured on that port.

1. Log on to the FastEthernet Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the brouter configuration from the switch
<brouter_port>	Specifies the port to configure as a brouter port. When modifying a brouter this is the port of the existing brouter to modify.
<vid>	Specifies the VLAN ID of the brouter. When creating a new brouter port, this is the VLAN ID assigned to the brouter port.

Variable	Value
<ip_address/mask>	Specifies the IP address and subnet mask of the router. When creating a new router, this is the IP address and subnet mask assigned. When modifying a router port, this is the new IP address and subnet mask to assign to the port. The subnet mask portion is expressed as a value between 0 and 32.
[routing enable]	Enables Layer 3 routing on the router port.

Displaying the router port configuration

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

Procedure steps

1. Log on to the User Exec mode in ACLI.
2. At the command prompt, enter the following command:

```
show router [port {router_port}] [vrf {vrfName}] [vrfs
{vrfid_list}]
```

Variable definitions

The following table describes the parameters for the **show router** command..

Variable	Value
port {router_port}	Narrows the scope of the command to the specified port(s). Omission of this parameter displays all ports.
vrf {vrfName}	Displays information for the VRF instance identified by the alphanumeric identifier (name)
vrfs {vrfid_list}	Displays information for the VRF instances listed by their numerical IDs. List values range from 0 to 3.

Job aid

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

Field	Description
Port	Specifies the router port number.
Router VID	Specifies the router VLAN ID.
IP/Mask	Specifies the IP address and subnet mask of the router.
IP Routing	Specifies whether IP routing is enabled or disabled on the router.

Chapter 11: Router port configuration using EDM

A router port is a single-port VLAN that can route IP packets and bridge all non-routable traffic. This chapter describes procedures you can use to configure a router port using Enterprise Device Manager (EDM).

Configuring a router port

Use the following procedure to configure and manage router 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.
4. In the work area, click the IP Address tab.
5. In the toolbar, click **Insert**.
6. Using the provided fields, create the new router port.
7. Click **Insert**.

Variable definitions

The following table describes the fields of the **IP Address** tab.

Field	Description
IpAddress	Specifies the IP address assigned to this router.
NetMask	Specifies the subnet mask associated with the router IP address.
VlanId	Specifies the VLAN ID associated with this router port.

Router port configuration using EDM

Field	Description
MacOffset	Specifies the MAC address offset associated with this router port.

Chapter 12: Static route configuration using ACLI

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

Configuring a static route

Use this procedure to configure 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

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Variable definitions

The following table describes the 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.
[disable]	Disables the specified static route.
[enable]	Enables 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

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip route static [<dest-ip>] [-s <subnet> <mask>] [vrf
{vrfName}] [vrfids {vrf_ids}]
```

Variable definitions

The following table describes the 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.
[vrf {vrfName}]	Specifies the VRF instance identified by the alphanumeric identifier (name) for which to display route information.
[vrfids {vrf_ids}]	Specifies the VRF instance(s) identified by the numerical ID(s) for which to display route information.

Job aid

The following table shows the field descriptions for the `show ip route static` 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 static routes, options are LOC (local route) or STAT (static route).
TYPE	Indicates the type of route as described by the Type Legend on the ACLI screen.
PRF	Specifies the route preference.

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 subnet. You can configure a maximum of 4 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

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Variable definitions

The following table describes the 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

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip mgmt route
```

Job aid

The following table shows the 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.
Status	<p>ACTIVE The status is ACTIVE when the following criteria are met:</p> <ul style="list-style-type: none"> • The management IP address is configured. • The management route next-hop resides in the same network as the management IP address. • The management VLAN is active—at least one member port is up. <p>INACTIVE The status is INACTIVE under all other circumstances.</p>

Displaying the multicast static IP routing table using ACLI

Use this procedure to display the multicast static IP routing table.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip mroute rpf <ipaddr> info
```

Creating a multicast-static IP routing table entry using ACLI

Use this procedure to enter an item into the multicast-static IP routing table.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip static-mroute create <ipaddr/mask> rpf <ip-addr>
[preference <value>]
```

Deleting a multicast-static IP routing table entry using ACLI

Use this procedure to delete an entry for a multicast-static IP routing table.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip static-mroute delete <ip-addr/mask> rpf <ip-addr>
```

Enabling a route in the multicast-static IP routing table using ACLI

Use this procedure to enable a route in the multicast-static IP routing table.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip static-mroute enable <ipaddr/mask> rpf <ip-addr>
```

Disabling a route in a multicast static IP routing table using ACLI

Use this procedure to disable a route in a multicast static IP routing table.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip static-mroute disable <ipaddr/mask> rpf <ip-addr>
```

Modifying the administrative distance of a multicast static IP route using ACLI

Use this procedure to modify the administrative distance of a multicast static IP route.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip static-mroute preference <value> <ipaddr/mask> rpf <ip-  
addr>
```


Chapter 13: Static route configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure static routes using Enterprise Device Manager (EDM).

Configuring static routes

Use the following procedure to configure static routes for the switch.

Prerequisites

- Enable IP routing globally.
- Enable IP routing, and then configure an IP address on the VLANs to use for routing.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IP**.
3. In the work area, select the **Static Routes** tab.
4. In the toolbar, click **Insert**.
5. Configure the parameters as required.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Static Routes** tab.

Field	Description
Dest	Specifies the destination IP address of the route. 0.0.0.0 is the default route.
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. The switch uses this value to choose the best route, the route with the smallest cost, to a certain destination. The range is 1 to 65535.
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 (inactive or active).

Displaying IP routes

Use the following procedure to display the different routes known to the switch.

The switch does not display routes until at least one port in the VLAN has a link.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IP**.
3. In the work area, select the **Routes** tab.

Variable definitions

The following table describes the fields of the **Routes** tab.

Field	Description
Dest	Specifies the destination address of the route.
Mask	Specifies the subnet mask that the route destination uses.
NextHop	Specifies the next hop in the listed route.

Field	Description
HopOrMetric	Specifies the Routing Information Protocol (RIP) hop count, Open Shortest Path First (OSPF) cost, or metric associated with the route.
Interface	Specifies the interface associated with the route.
Proto	Specifies the protocol associated with the route.
PathType	Specifies the route path type: <ul style="list-style-type: none"> • i: indirect • d: direct • A: alternative • B: best • E: ECMP • U: unresolved
Pref	Specifies the preference value associated with the route.

Filtering route information

Use the following procedure to filter the **Routes** tab to display only the desired switch routes.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IP**.
3. In the work area, select the **Routes** tab.
4. In the toolbar, click **Filter**.

Variable definitions

The following table describes the fields of the **Filter** tab.

Field	Description
Condition	When you use multiple filter expressions on the tab, the condition joins them together.
Ignore Case	Denotes whether filters are case sensitive or insensitive.
Column	Denotes the type of criteria that will apply to values for filtering.

Field	Description
All Records	Clears the filters, and displays all rows.
Dest	Select this check box and enter a value to filter on the route destination value.
Mask	Select this check box and enter a value to filter on the route destination subnet mask value.
NextHop	Select this check box and enter a value to filter on the next hop value of the route.
HopOrMetric	Select this check box and enter a value to filter on the hop count or metric of the route.
Interface	Select this check box and enter a value to filter on the associated interface of the route.
Proto	Select this check box and enter a value to filter on the route protocol.
PathType	Select this check box and enter a value to filter on the route path type.
Pref	Select this check box and enter a value to filter on the route preference value.

Creating a multicast-static IP routing table entry

Use this procedure to create an entry in the multicast-static IP routing table.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Static MRoutes** tab.
4. On the toolbar, click **Insert**.
5. Configure the parameters as required.
6. Click **Insert**.

Variable definitions

Variable	Value
IpAddressType	Specifies the type of IP Address (ipv4).

Variable	Value
IpAddress	Specifies the IP address of the destination network.
Mask	Specifies the mask of the destination network.
RpfAddressType	Specifies the type of address for the reverse path forwarding address (ipv4).
RpfAddress	Specifies the reverse path forwarding address.
Preference	Specifies the administrative distance of the static multicast route

Displaying TCP information for the switch

Use the following procedure to display Transmission Control Protocol (TCP) information for the switch.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **TCP/UDP**.

Variable definitions

The following table describes the fields of the **TCP Globals** tab.

Field	Description
RtoAlgorithm	Specifies the algorithm 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.

Displaying TCP connections

Use the following procedure to display information on the current TCP connections the switch maintains.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **TCP/UDP**.
3. In the work area, click the **TCP Connections** tab.

Variable definitions

The following table describes the fields of the **TCP Connections** tab.

Field	Description
LocalAddressType	Specifies the local IP address type for this TCP connection.
LocalAddress	Specifies the local IP address for this TCP connection. In the case of a connection in the listen state, which is willing to accept connections for any IP interface associated with the node, the value is 0.0.0.0.
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.

Displaying TCP listeners

Use the following procedure to display information on the current TCP listeners on the switch.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **TCP/UDP**.
3. In the work area, click the **TCP Listeners** tab.

Variable definitions

The following table describes the fields of the **TCP Listeners** tab.

Field	Description
LocalAddressType	Specifies the IP address type of the local TCP listener.
LocalAddress	Specifies the local IP address of the TCP listener. The switch represents the value of this field in three possible ways, depending on the characteristics of the listening application: <ol style="list-style-type: none"> 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

Displaying UDP endpoints

Use the following procedure to display information on the UDP endpoints currently maintained by the switch.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **TCP/UDP**.
3. In the work area, click the **UDP Endpoints** tab.
4. In the toolbar, click **Refresh** to refresh the information.

Variable definitions

The following table describes the fields of the **UDP Listeners** tab.

Field	Description
LocalAddressType	Specifies the local address type (IPv6 or IPv4).
LocalAddress	<p>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 is 0.0.0.0.</p> <ol style="list-style-type: none"> 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 the switch accepts datagrams from all remote systems, this value is a zero-length octet string. Otherwise, the address of the remote system from which the switch accepts datagrams (or to which the switch sends all datagrams) is displayed with the RemoteAddressType identifying the supported address type.
RemotePort	Displays the remote port number. If the switch accepts datagrams from all remote systems, this value is zero.

Field	Description
Instance	Distinguishes between multiple processes connected to the same UDP endpoint.
Process	Displays the ID for the UDP process.

Chapter 14: 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.

Configuring the OSPF hardware mode

Use this procedure to configure the OSPF hardware-compatibility mode.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip ospf op-mode {5510 | non-5510}
```

Variable definitions

The following table describes the parameters for the `ip ospf op-mode` command.

Variable	Value
5510	Indicates when 5510 switches participate in OSPF. This includes 5510 to 5510, and 5510 to 5600.
non-5510	Indicates 5510 switch does not participate in OSPF. This configuration is required when 5500 (except 5510) and 5600 are connected to 5500 (except 5510), 5600, and 8600.

Configuring the router ID

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

Procedure steps

1. Log on to the Router Configuration mode in ACLI.
2. At the command prompt, enter 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 global OSPF status

Use this procedure to configure the status of OSPF globally on the switch.

By default, OSPF is disabled.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter 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 global OSPF parameters

Use this procedure to define the global OSPF parameters, including default cost metric, RFC 1583 compatibility, OSPF holddown timer, and OSPF system traps.

Procedure steps

1. Log on to the Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] default-cost {ethernet | fast-ethernet |  
gig-ethernet | ten-gig-ethernet} <metric_value>  
rfc1583-compatibility enable
```

```
timers basic holddown <timer_value>
trap enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets the specified parameter to the default value.
[no]	Disables the specified feature (applicable only to rfc-1583-compatibility enable and trap enable).
{ethernet fast-ethernet gig-ethernet ten-gig-ethernet} <metric_value>	Specifies the default cost metric to assign to the specified port type. The metric value is an integer between 1 and 65535. The default values are as follows: <ul style="list-style-type: none"> • ethernet (10 Mb/s): 100 • fast-ethernet (100 Mb/s): 10 • gig-ethernet (1000 Mb/s): 1 • ten-gig-ethernet (10000 Mb/s): 1
rfc1583-compatibility enable	Enables RFC 1583 compatibility on the switch.
timers basic holddown <timer_value>	Specifies a holddown timer value between 3 and 60 seconds.
trap enable	Enables OSPF system traps.

Displaying global OSPF parameters

Use this procedure to display global OSPF parameters.

Procedure steps

1. Log on to the User EXEC Configuration mode in ACLI
2. At the command prompt, enter the following command:

```
show ip ospf
```

Configuring OSPF area parameters

Use this procedure to configure OSPF area parameters.

Procedure steps

1. Log on to the Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] area <area-id>
[default-cost {0-16777215}]
[import {external | noexternal | nssa}]
[import-summaries {enable}]
[range {subnet_mask} [{nssa-entlink | summary-link}]
[advertise-mode {no-summarize | summarize | suppress}]
[advertise-metric {0-65535}]
```

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: <ul style="list-style-type: none"> • external: specifies a normal area • noexternal: specifies a stub area • nssa: specifies an NSSA <p>* Note: The configuration of a totally stubby area (no summary advertising) is a two step process. First, define an area with</p>

Variable	Value
	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 {subnet_mask} [{nssa-entlink summary-link}] [advertise-mode {no-summarize summarize suppress}] [advertise-metric {0-65535}]	Specifies range parameters for the OSPF area.

Displaying OSPF area configuration

Use this procedure to display OSPF area configuration.

Procedure steps

1. Log on to the User EXEC Configuration mode in ACLI.
2. At the command prompt, enter 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.

Procedure steps

1. Log on to the User EXEC Configuration mode in ACLI.
2. At the command prompt, enter 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 routing on an interface

Use this procedure to enable OSPF routing on an interface and assign the interface to an OSPF area.

Procedure steps from Router Configuration mode

1. Log on to the Router Configuration mode or Interface Configuration mode in ACLI.
2. At the command prompt (Router Configuration mode), enter the following command:

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

OR

At the command prompt (Interface Configuration mode), enter the following command:

```
[no] ip ospf vlan <vid> 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).
<vid>	Specifies the VLAN ID to be enabled for OSPF routing.

Assigning an interface to an OSPF area

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

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter 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 the OSPF properties for an interface

Use this procedure to configure OSPF properties for an interface.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip ospf
advertise-when-down enable
cost <interface_cost>
dead-interval <interval>
hello-interval <interval>
mtu-ignore enable
network {broadcast | passive}
priority <0-255>
retransmit-interval <interval>
transit-delay <interval>
```

Variable definitions

The following table describes the command variables.

Variable	Value
advertise-when-down enable	Enables advertisement of the OSPF interface even when the interface is operationally unavailable.
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.
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.

Variable	Value
priority <priority_value>	Assigns a priority to the interface for the purposes of Designated Router election. This is an integer value between 0 and 255.
retransmit-interval <interval>	Defines the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. This is an integer value between 0 and 3600.
transit-delay <interval>	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 0 and 3600.

Displaying OSPF interface timers

Use this procedure to display OSPF interface timers

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip ospf timer {interface [FastEthernet <portlist> | vlan
<vid>] | virtual-links}
```

OR

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

Variable definitions

The following table describes the command variables.

Variable	Value
interface [FastEthernet <portlist> vlan <vid>]	Displays configured timers for the specified interface. If no interface is specified, all interface timers are displayed.
virtual-links	Displays configured timers for virtual links.

Displaying OSPF interface configurations

Use this procedure to display OSPF interface configurations.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip ospf interface [FastEthernet <portlist> | vlan <vid>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[FastEthernet <portlist> 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.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter 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.

Procedure steps

1. Log on to the Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Variable definitions

The following table describes the command variables.

Variable	Value
no	Removes the ASBR configuration for the router.

Configuring the OSPF authentication type for an interface

Use this procedure to configure the interface authentication type.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter 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.

Defining simple authentication keys for OSPF interfaces

Use this procedure to configure an interface authentication password.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip ospf 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.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter 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.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip ospf authentication {interface [FastEthernet
<portlist> | vlan <vid>] | virtual-links}
```

Variable definitions

The following table describes the command variables.

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

Procedure steps

To specify the primary MD5 key, enter the following from the Interface Configuration command mode:

```
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.

Procedure steps

To display OSPF authentication configuration for interfaces, enter the following from the User EXEC command mode:

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

Configuring a virtual link

Use this procedure to create a virtual interface.

Procedure steps

To create a virtual interface, enter the following from the OSPF Router Configuration command mode:

```
[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-i
nterval <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.
<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.

Variable	Value
authentication-key <WORD>	Specifies the unique identifier assigned to the authentication key.
authentication-type	Specifies one of the following authentication types: <ul style="list-style-type: none"> • 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.

Procedure steps

To create a virtual interface message digest key, enter the following from the OSPF Router Configuration command mode:

```
area virtual-link message-digest-key <A.B.C.D.>  
<A.B.C.D./0-32> <1-255> md5-key <WORD>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a virtual interface message digest key.
<A.B.C.D.>	Specifies the transit area Id expressed as an IP address.
<W.X.Y.Z>	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.

Configuring automatic virtual links

Use this command to enable global automatic Virtual Link creation.

Procedure steps

1. Log on to the Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

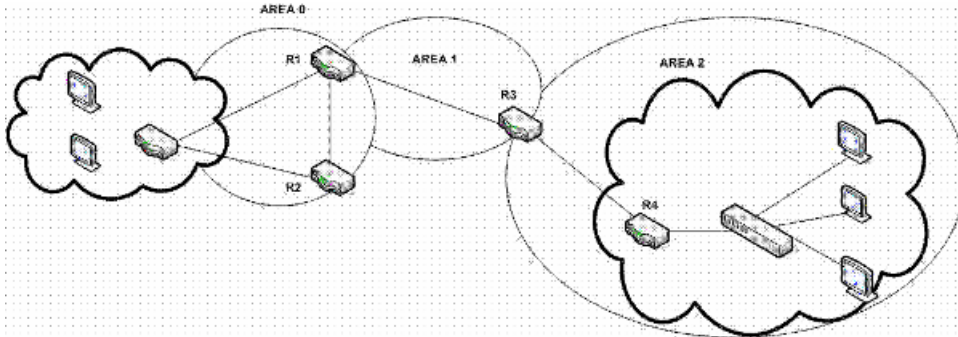
Variable definitions

The following table describes the command variables.

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

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 9: 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)
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 10: 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.

Procedure steps

1. Log on to the User EXEC Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Displaying OSPF virtual neighbors

Use this procedure to display OSPF virtual neighbors.

Procedure steps

1. Log on to the User EXEC Configuration mode in ACLI.
2. At the command prompt, enter the following command:

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

Configuring an OSPF host route

Use this procedure to add a host to a router.

Procedure steps

To add a host to a router, enter the following from the OSPF Router Configuration command mode:

```
[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.
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.

Procedure steps

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


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

Displaying the OSPF link state database

Use this procedure to display the OSPF link state database.

Procedure steps

To display the OSPF link state database, enter the following from the User EXEC command mode:

```
show ip ospf lsd
{[area <area-id>]
[lsa-type <type>]
[lsid <ip_address>]
[adv-rtr <router_id>] |
detail [<router_id>]}
```

Variable definitions

The following table describes the command variables.

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

Displaying the external link state database

Use this procedure to display the external link state database.

Procedure steps

To display OSPF ASE LSAs, enter the following from the User EXEC command mode:

```
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

Procedure steps

To immediately initiate an SPF run, enter the following from the Global Configuration command mode:

```
ip ospf spf-run
```

Displaying OSPF default port metrics

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

Procedure steps

To display OSPF default metrics, enter the following from the User EXEC command mode:

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

Displaying OSPF statistics

Use this procedure to display OSPF statistics.

To clear OSPF statistics counters, use the `clear ip ospf counters` command.

Procedure steps

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

```
show ip ospf stats
```

Displaying OSPF interface statistics

Use this procedure to display OSPF interface statistics.

Procedure steps

To display OSPF interface statistics, enter the following from the User EXEC command mode:

```
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.

Procedure steps

To clear OSPF statistics counters, enter the following from the Global Configuration command mode:

```
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 15: OSPF configuration examples using ACLI

The following sections provide Open Shortest Path First (OSPF) configuration examples using ACLI.

Basic OSPF configuration examples

This section contains the steps necessary for the initial configuration of OSPF on the switch. For more advanced configuration examples, see [Advanced OSPF configuration examples](#) on page 216.

*** Note:**

Many of the following configuration examples use a brouter port to create a connection to the network core. This practice does not imply that a brouter port is the only means through which you can establish a core connection. The use of a brouter port is only one of many ways to create such a connection.

Basic OSPF configuration

A basic OSPF configuration learns OSPF routes from other OSPF devices, and propagate routes to other OSPF devices. The following procedure provides the steps to create a basic OSPF configuration:

1. Log into User EXEC mode.
`5530-24TFD> enable`
2. Log into Global Configuration mode.
`5530-24TFD# config terminal`
3. Enable IP routing globally.
`5530-24TFD(config)# ip routing`
4. Enable OSPF globally.
`5530-24TFD(config)# router ospf en`

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

```
5530-24TFD(config)# router ospf
```

*** Note:**

The remainder of this procedure refers to VLAN 35. Although this example uses VLAN 35, you can use any port type VLAN.

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

```
5530-24TFD(config)# vlan create 35 type port 1
```

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

```
5530-24TFD(config)# interface vlan 35
```

8. Enable IP routing on VLAN 35.

```
5530-24TFD(config-if)# ip routing
```

9. Assign an IP address to VLAN 35.

```
5530-24TFD(config-if)# ip address 1.1.2.25 255.255.255.0
```

10. Enable OSPF on VLAN 35.

```
5530-24TFD(config-if)# ip ospf en
```

11. Return to Global Configuration mode.

```
5530-24TFD(config-if)# exit
```

12. By default all ports belong to a newly created VLAN. The following command removes all ports from VLAN 35.

```
5530-24TFD(config)# vlan members remove 35 all
```

13. Add ports 1 through 10 to VLAN 35.

```
5530-24TFD(config)# vlan members add 35 1-10
```

Basic ASBR configuration

OSPF uses the Autonomous System Boundary Router (ASBR) to import routes that come from non-OSPF sources, such as the following:

- Local interfaces that are not part of OSPF.
- Routing Information Protocol (RIP) interfaces.
- RIP learned routes.
- Static routes.

Use this quick reference to help configure OSPF to import these types of routes and allow the rest of the OSPF network to learn them as OSPF routes. To create a basic ASBR configuration, perform the following procedure:

1. Log into User EXEC mode.

```
5530-24TFD> enable
```

2. Log into Global Configuration mode.

```
5530-24TFD# config terminal
```

3. Log into the OSPF router configuration mode.

```
5530-24TFD(config)# router ospf
```

4. Enable ASBR functionality.

```
5530-24TFD(config-router)# as-boundary-router en
```

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

```
5530-24TFD(config-router)# redistribute rip en
5530-24TFD(config-router)# redistribute direct en
5530-24TFD(config-router)# redistribute static en
```

6. Return to Global Configuration mode.

```
5530-24TFD(config-router)# exit
```

7. After you use the commands in step 5 to select the types of routes to redistribute, apply the changes globally with the following commands.

```
5530-24TFD(config)#ip ospf apply redistribute rip
5530-24TFD(config)#ip ospf apply redistribute direct
5530-24TFD(config)#ip ospf apply redistribute static
```

Configuring ECMP for OSPF

To configure Equal Cost Multipath (ECMP) with OSPF, use the following procedure.

The 5510 models do not support ECMP with OSPF.

1. Log into User EXEC mode.

```
5530-24TFD> enable
```

2. Log into Global Configuration mode.

```
5530-24TFD# config terminal
```

3. Set the number of ECMP paths to use with OSPF. OSPF can use up to four ECMP paths.

```
5530-24TFD(config)# ospf maximum-path 2
```

This command tells the router to use up to two equal-cost paths to get to any OSPF network destination.

4. Verify the configuration.

```
5530-24TFD(config)# show ecmp
```

Advanced OSPF configuration examples

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

The Avaya Ethernet Routing Switch 5000 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 router port or on a VLAN. The following section demonstrates the creation of the example OSPF interface illustrated in the following figure.




Figure 33: OSPF interface example topology

To create the OSPF interface illustrated for router R1, perform the following procedure:

1. Configure brouter port OSPF interface.

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

```
5530-24TFD# config terminal
5530-24TFD(config)# interface fast 2
5530-24TFD(config-if)# brouter port 2 vlan 2134 subnet
10.1.1.21/30
5530-24TFD(config-if)# router ospf
5530-24TFD(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 then enable OSPF on this interface.

```
5530-24TFD(config)# vlan create 2 type port
5530-24TFD(config)# spanning-tree stp 1 add-vlan 2
5530-24TFD(config)# vlan member add 2 1
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 172.3.1.1
255.255.255.0
5530-24TFD(config-if)# router ospf
5530-24TFD(config-router)# network 172.3.1.1
```

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

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# router-id 1.1.1.1
5530-24TFD(config-router)# exit
5530-24TFD(config)# router ospf enable
```

OSPF security configuration example using Message Digest 5

In the configuration example illustrated in the following figure, routers R1 and R2 use MD5 authentication.




Figure 34: MD5 configuration example

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

1. Configure MD5 authentication on R1.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip ospf message-digest-key 2 md5
qwsdf89
5530-24TFD(config-if)# ip ospf primary-md5-key 2
```

```
5530-24TFD(config-if)# ip ospf authentication-type
message-digest
```

2. Configure MD5 authentication on R2.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip ospf message-digest-key 2 md5
qw sdf89
5530-24TFD(config-if)# ip ospf primary-md5-key 2
5530-24TFD(config-if)# ip ospf authentication-type
message-digest
```

Configuring OSPF network types

Use OSPF network types to allow OSPF-neighboring between routers over different types of network infrastructures. With this feature, you can configure each interface to support the various network types.

In the following figure, VLAN 2 on Avaya Ethernet Routing Switch 5000 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 35: OSPF network example

To create this configuration for router R1, use the following commands:

```
5530-24TFD(config)# vlan create 2 type port
5530-24TFD(config)# vlan mem add 2 1
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 172.3.1.1 255.255.255.0
5530-24TFD(config-if)# ip ospf network passive
```

The Avaya Ethernet Routing Switch 5000 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 the interface network to be included in OSPF without generating LSAs or forming adjacencies. This interface type is typically used on an access network. This

type 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 5000 Series ; no user intervention is required. The Avaya Ethernet Routing Switch 5000 Series automatically becomes an OSPF ABR when it has operational OSPF interfaces that belong to more than one area.

In the following figure, the Avaya Ethernet Routing Switch 5000 Series R1 automatically becomes an OSPF ABR after you configure it with an OSPF interface for area 0.0.0.0 and 0.0.0.2.

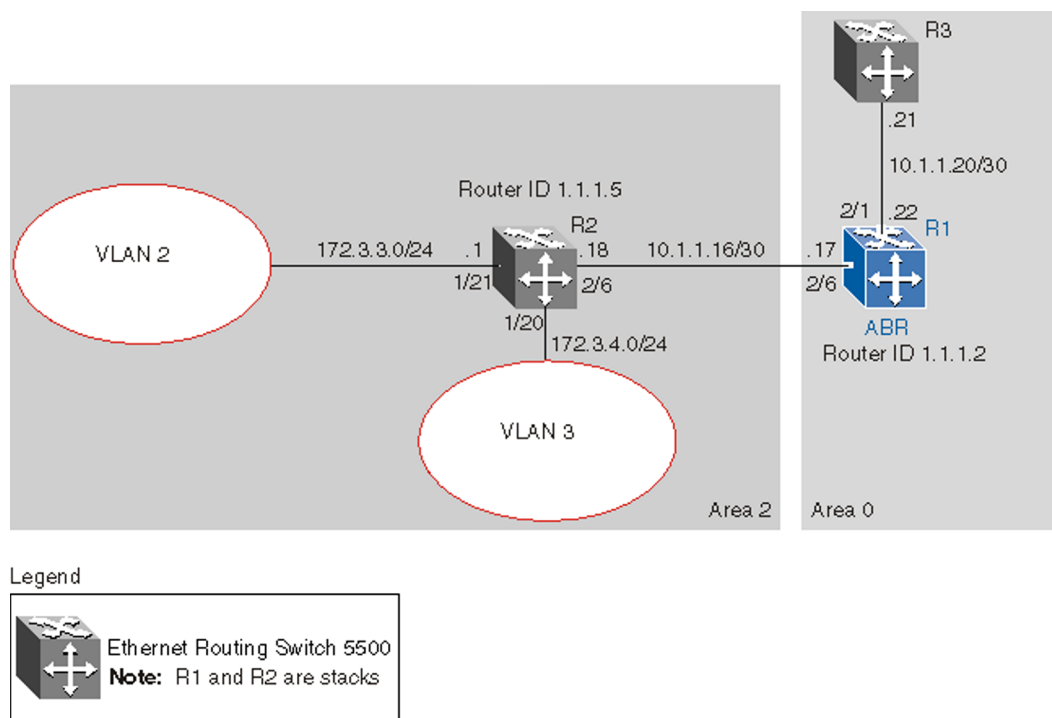


Figure 36: ABR configuration example

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.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100 subnet
10.1.1.17/30
```

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.

```
5530-24TFD(config)# interface fast 2/1
5530-24TFD(config-if)# brouter port 2/1 vlan 200 subnet
10.1.1.22/30
5530-24TFD(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.

```
5530-24TFD(config-router)# router-id 1.1.1.2
5530-24TFD(config-router)# area 0.0.0.2
5530-24TFD(config-router)# network 10.1.1.17 area
0.0.0.2
5530-24TFD(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 a metric of 100.

```
5530-24TFD(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 ASBRs

An Autonomous System Border Router (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 5000 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 shows an Avaya Ethernet Routing Switch 5000 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 37: 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 then enabling RIP on this interface.

```
5530-24TFD(config)# interface fast 1/31
5530-24TFD(config-if)# brouter port 1/31 vlan 100 subnet
10.1.1.41/30
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 10.1.1.41
```

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

```
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 100
5530-24TFD(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 then enable OSPF on this interface.

```
5530-24TFD(config)# interface fast 2/7
5530-24TFD(config-if)# brouter port 2/7 vlan 200 subnet
10.1.1.2/30
5530-24TFD(config-if)# router ospf
5530-24TFD(config-router)# network 10.1.1.2
```

4. Make R1 the ASBR.

Configure R1 as an ASBR, and then assign the OSPF router ID.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# as-boundary-router enable
5530-24TFD(config-router)# router-id 1.1.1.3
5530-24TFD(config)# router ospf enable
```

5. Configure OSPF route distribution.

Configure OSPF route distribution to import RIP into OSPF. The Avaya Ethernet Routing Switch 5000 Series distributes the RIP routes as AS-external LSA (LSA type 5), using external metric type 1.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# redistribute rip enable
metric 10 metric-type type1
5530-24TFD(config)# ip ospf apply redistribute rip
```

6. Configure a route policy.

You must use a route policy 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.

```
5530-24TFD(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 created in step 6 to RIP interface 10.1.1.41.

```
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip out-policy allow
```

The configuration steps in the preceding example distribute all OSPF routes to RIP. However, there are times when it is 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.

```
5530-24TFD(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 the match-protocol value. This configuration causes the default route to be advertised through RIP only if OSPF is operational.

```
5530-24TFD(config)# route-map Policy_Default permit 1
enable match protocol ospf set injectlist default
5530-24TFD(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.

```
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip out-policy Policy_Default
```

Stub area configuration example

In the following configuration example, the Avaya Ethernet Routing Switch 5000 Series R1 is configured in stub area 2, and R2 is configured as a stub ABR for area 2.

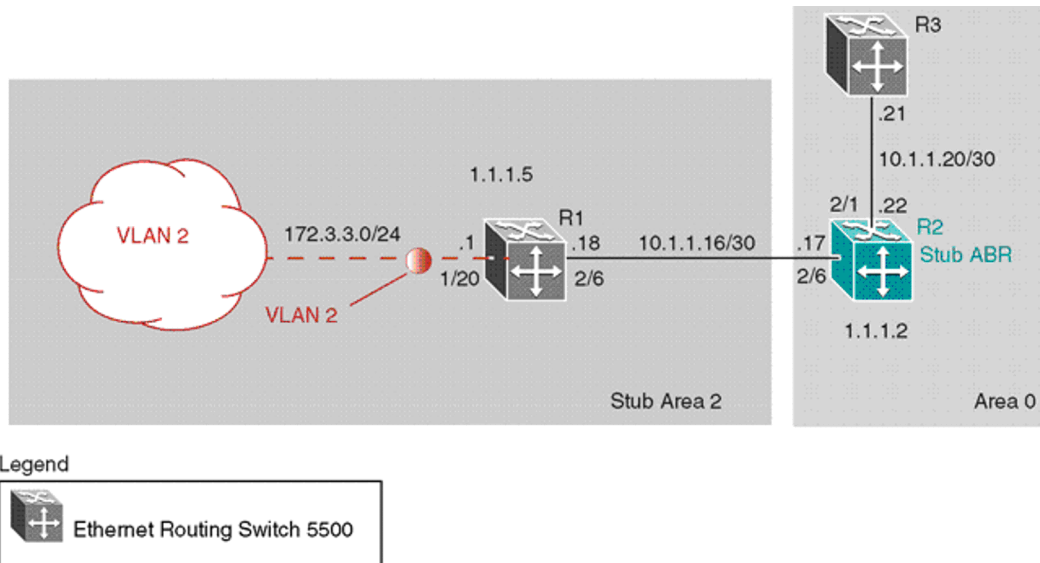


Figure 38: OSPF stub area example

*** 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 procedure to perform the stub area configuration illustrated in this example:

*** Note:**

This example assumes that you have enabled global IP routing on the switch. Enable global IP routing on the switch in Global Configuration mode using the `ip routing` command.

1. Configure router R1.
 - a. Configure the OSPF interface on R1.
Configure port 2/6 as a brouter port in VLAN 100.


```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100
5530-24TFD(config-if)# subnet 10.1.1.18/30
```
 - b. Configure VLAN 2 on R1.
Create VLAN 2 and assign an IP address to it.


```
5530-24TFD(config)# vlan create 2 type port
5530-24TFD(config)# vlan mem add 2 1/20
```



```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 172.3.3.1
255.255.255.0
```

c. 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 then enable OSPF on these interfaces.

```
5530-24TFD(config-router)# router-id 1.1.1.5
5530-24TFD(config-router)# area 0.0.0.2 import
noexternal
5530-24TFD(config-router)# network 10.1.1.18 area
0.0.0.2
5530-24TFD(config-router)# network 172.3.3.1 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

2. Configure router R2.

a. Configure the OSPF interface on R2.

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

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100
subnet 10.1.1.17/30
```

b. Configure the second OSPF interface on R2.

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

```
5530-24TFD(config)# interface fast 2/1
5530-24TFD(config-if)# brouter port 2/1 vlan 300
subnet 10.1.1.22/30
5530-24TFD(config-if)# ip ospf enable
```

c. 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. R2 will originate only summary LSA for default route into area 2. Note that, by default, OSPF interface 10.1.1.22 is in OSPF area 0.0.0.0. Because you add one additional area of 0.0.0.2, and OSPF interface 10.1.1.17 is added to area 0.0.0.2, R2 automatically becomes a stub ABR.

```
5530-24TFD(config-router)# router-id 1.1.1.2
5530-24TFD(config-router)# area 0.0.0.2 import
noexternal
5530-24TFD(config-router)# no area 0.0.0.2
import-summary enable
5530-24TFD(config-router)# area 0.0.0.2 default-cost
10
5530-24TFD(config-router)# network 10.1.1.17 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

NSSA configuration example

The following NSSA configuration example demonstrates an Avaya Ethernet Routing Switch 5000 Series configured as a NSSA ASBR router.

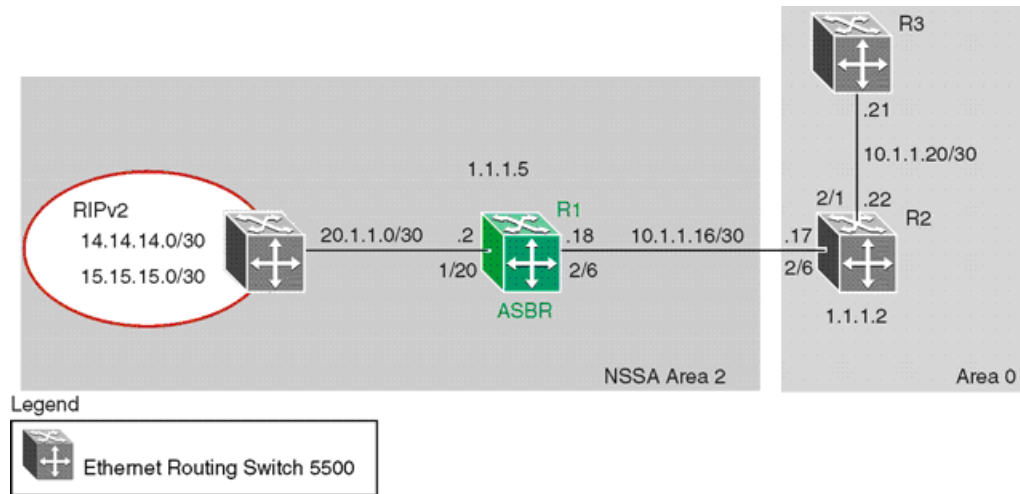


Figure 39: NSSA configuration example

To configure the example illustrated above, perform the following procedure:

Configure router R1.

- a. Configure the RIP interface on R1.

Configure port 1/20 as a brouter port in VLAN 100, and then enable RIP on this interface.

```
5530-24TFD(config)# interface fast 1/20
5530-24TFD(config-if)# brouter port 1/20 vlan 100
subnet 20.1.1.2/30
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 20.1.1.2
```

- b. Enable RIP globally, and then configure the RIP version 2 interface.

```
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip receive version rip2
send version rip2
```

- c. Configure the OSPF interface on R1.

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

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 200
subnet 10.1.1.18/30
```

- d. 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 then enable OSPF on the interface.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# as-boundary-router
enable 5530-24TFD(config-router)# router-id 1.1.1.5
5530-24TFD(config-router)# area 0.0.0.2 import nssa
5530-24TFD(config-router)# network 10.1.1.18 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

- e. Configure a route policy to distribute direct and OSPF routes to RIP.

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

```
5530-24TFD(config)# route-map Rip_Dist permit 1
enable match protocol direct,ospf set metric-type
type1
```

- f. Apply the Rip_Dist route policy to RIP out policy.

```
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip out-policy Rip_Dist
```

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

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# redistribute rip enable
metric-type type1
5530-24TFD(config)# ip ospf apply redistribute rip
```

Controlling NSSA external router advertisements

In an OSPF NSSA, the ABR uses the NSSA N/P-bit (in the OSPF hello packets options field) to advertise external routes to other areas. If the NSSA N/P-bit is true, the ABR exports the external route. True is the default setting for the Avaya Ethernet Routing Switch 5000 Series. When the NSSA N/P-bit is not set true, the ABR drops the external route. You can create a route policy on the Avaya Ethernet Routing Switch 5000 Series to manipulate the N/P-bit value.

For example, the following illustration 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.

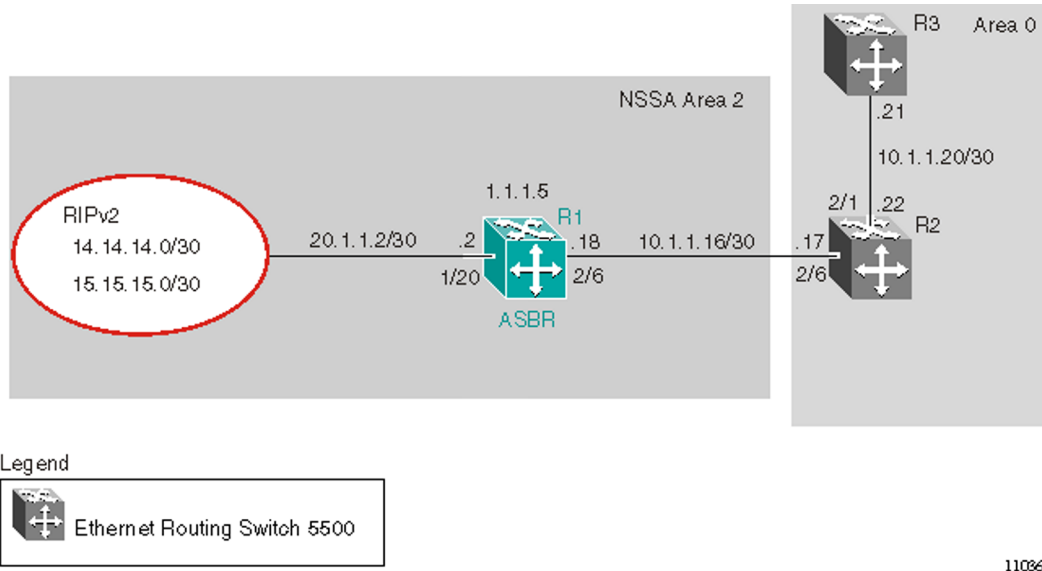


Figure 40: External router advertisement example

The following procedure provides the commands to replicate this configuration example:

1. Configure the RIP interface.

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

```
5530-24TFD(config)# interface fast 1/20
5530-24TFD(config-if)# brouter port 1/20 vlan 200 subnet
20.1.1.2/30
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 20.1.1.2
```

2. Globally enable RIP, and then configure a RIP interface for RIP version 2.

```
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 200
5530-24TFD(config-if)# ip rip receive version rip2 send
version rip2
```

3. Configure the OSPF interface.

Configure port 2/6 as a brouter port.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(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.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# router-id 1.1.1.5
5530-24TFD(config-router)# as-boundary-router enable
5530-24TFD(config-router)# area 0.0.0.2 import nssa
```

```
5530-24TFD(config-router)# network 10.1.1.18 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

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

```
5530-24TFD(config)# route-map Rip_Dist permit 1 enable
match protocol direct,ospf set metric-type type1
```

6. Apply route policy to RIP out policy.

```
5530-24TFD(config)# interface vlan 200
5530-24TFD(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.

```
5530-24TFD(config)# ip prefix-list 15net 15.15.15.0/24
5530-24TFD(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.

```
5530-24TFD(config)# route-map P_bit permit 1 enable
match network 15net set nssa-pbit enable
5530-24TFD(config)# route-map P_bit permit 2 enable
match network 14net
5530-24TFD(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.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# redistribute rip enable
metric-type type1 route-policy P_bit
5530-24TFD(config)# ip ospf apply redistribute rip
```

Configuring a multi-area complex

The multi-area complex configuration example in this section uses five Avaya Ethernet Routing Switch 5000 Series devices (R1 to R5) in a multi-area configuration.

Previous sections of this chapter describe many of the concepts and topology descriptions in this example configuration. The concepts shown in those examples are combined in this example configuration to show a real world topology example with command descriptions.

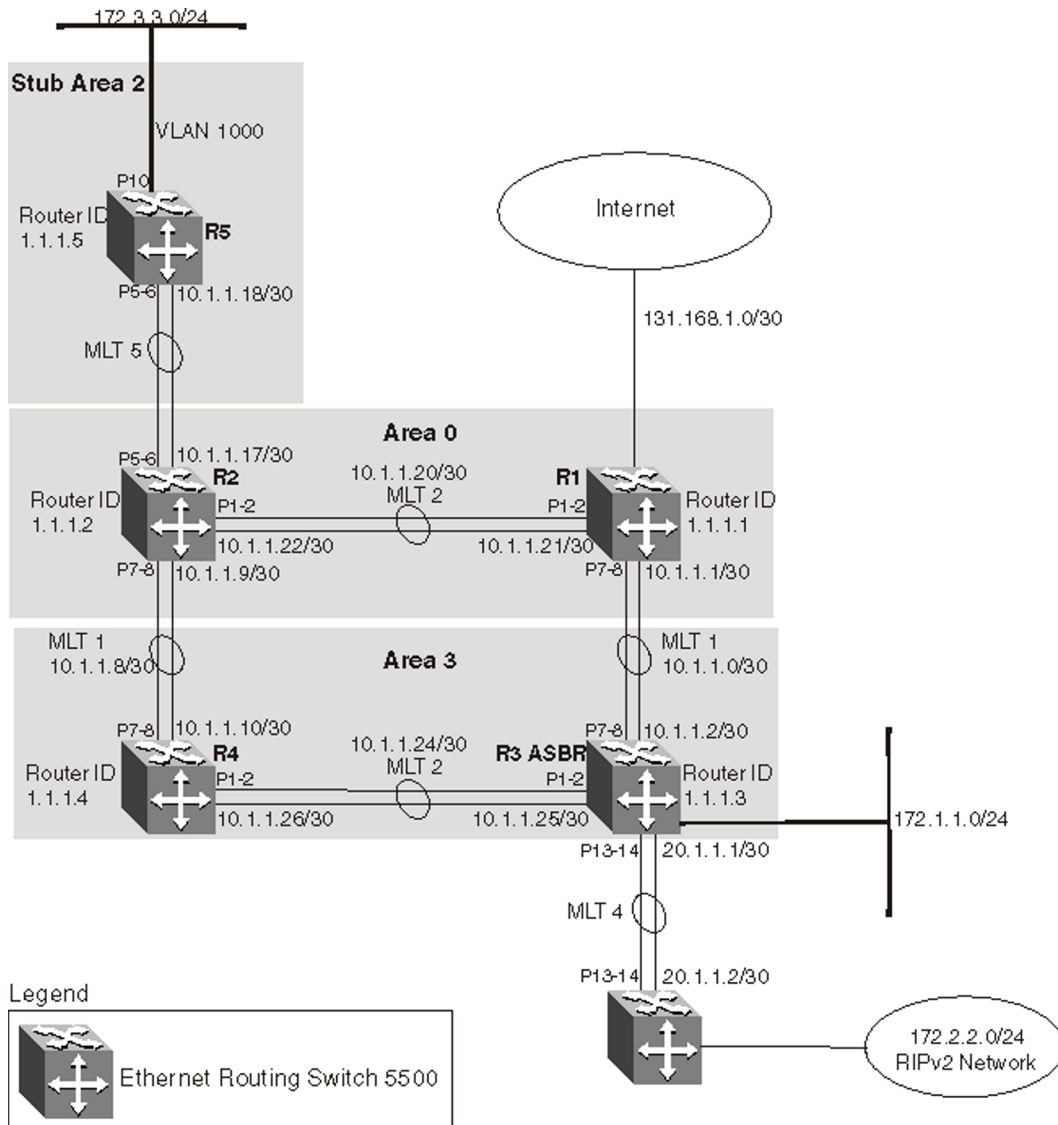


Figure 41: Multi-area complex example

For this configuration example, the Avaya Ethernet Routing Switch 5000 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 distributes OSPF routes to RIP and RIP routes to OSPF.
- R4 is an OSPF internal router in area 3.
- R5 is an internal OSPF stub router in area 2.
- All interfaces are Ethernet, therefore the OSPF interfaces are broadcast.

- The interface priority value on R5 is 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 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 provides the commands 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 that matches 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
```

OSPF configuration examples using ACLI

```
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
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 FastEthernet 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
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-
unregister
ed-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-
unregister

```

```

ed-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
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 FastEthernet 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
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
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 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 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-
unregister
ed-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-
unregister
ed-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
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 FastEthernet 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 FastEthernet ALL
spanning-tree port 10 learning disable
exit
interface FastEthernet 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
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
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 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
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

```

OSPF configuration examples using ACLI

```
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
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 FastEthernet 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
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

```

OSPF configuration examples using ACLI

```
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
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 FastEthernet ALL
spanning-tree port 5-6,24-26 learning 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 FastEthernet 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
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
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 10.1.1.21 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
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
```

OSPF configuration examples using ACLI

```
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 Poison Proxy
-----
10.1.1.2 1 false false false false true true false false
10.1.1.25 1 false false false false true truefalse 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

```

OSPF configuration examples using ACLI

```
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
-----
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
-----
```

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
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

```

OSPF configuration examples using ACLI

```
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
10.1.1.20.255.255.252 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
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
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

```

OSPF configuration examples using ACLI

```
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 172.3.3.1255.255.255.252 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
-----
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 to 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 11: OSPF neighbor states

Step	State	Description
1	Down	Indicates that a neighbor was configured manually, but the router did not receive 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 5000 Series does not support the 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 and 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 LSAs for missing or corrupt routes.
8	Full	Indicates the normal full adjacency state.

OSPF neighbor state information

You can access neighbor state information by using the `show ip ospf neighbor` command.

```
5530-24TFD#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. Several possible causes exist 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 use a high MTU size, and the other router a smaller value. Depending on the size of the LSDB, 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 5000 Series automatically checks for OSPF MTU mismatches.

In the Avaya Ethernet Routing Switch 5000 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 16: OSPF configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure the Open Shortest Path First (OSPF) protocol using Enterprise Device Manager (EDM).

The 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.

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.

Configuring Global OSPF properties

Use the following procedure to configure global OSPF parameters.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the **General** tab, configure the parameters as required.
4. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **General** tab.

Field	Description
RouterId	Specifies the unique ID of the router in the Autonomous System.
AdminStat	Specifies the administrative status 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 checksums of the external link state advertisements contained in the link state database. This sum determines if the link state database of the router changes, and compares the link state databases of two routers.
OriginateNewLsas	Specifies the number of new link state advertisements that the router originates. This number increments 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.
10MbpsPortDefaultMetric	Specifies the default metric of a 10 Mbps port. This is an integer value between 1 and 65535. Default value is 100.

Field	Description
100MbpsPortDefaultMetric	Specifies the default metric of a 100 Mbps port. This is an integer value between 1 and 65535. Default value is 10.
1000MbpsPortDefaultMetric	Specifies the default metric of a 1000 Mbps port. This is an integer value between 1 and 65535. Default value is 1.
10000MbpsPortDefault Metric	Specifies the default metric of a 10000 Mbps port. This is an integer value between 1 and 65535. Default value is 1.
TrapEnable	Specifies whether OSPF traps are enabled. The default setting is disabled.
AutoVirtLinkEnable	Specifies the status of OSPF automatic Virtual Link creation. The default setting is disabled.
SpfHoldDownTime	Specifies the SPF hold down timer value, which is an integer between 3 and 60. The 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 that advertise the same destination. When you enable this field, the preference rule will be the same as specified by RFC 1583. When disabled, the new preference rule, as described in RFC 2328, will apply. 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 the last SPF calculation was done.

Configuring an automatic virtual link

Automatic virtual links require more system resources than manually-configured virtual links. Automatic virtual links are removed when the transit area is deleted or when the router is no longer an ABR.

Use the following procedure to configure an automatic virtual link to provide an automatic, dynamic backup link for vital OSPF traffic.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **General** tab.
4. Type the router ID for one of the end point ABRs in the **RouterId** field .
5. Select the **AutoVirtLinkEnable** check box.
6. In the toolbar, click **Apply**.
7. On the remote ABR to use for the virtual link, repeat the preceding steps.

Configuring an OSPF area

Use the following procedure to configure an OSPF area.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Areas** tab.
4. In the toolbar, click **Insert**.
5. Using the fields provided, configure the OSPF area.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Areas** tab.

Field	Description
Areald	Specifies the unique identifier for the area. Area ID <i>0.0.0.0</i> is the OSPF backbone.
ImportAsExtern	Specifies the area type by defining its support for importing Autonomous System external link state advertisements. The options available are: <ul style="list-style-type: none"> • 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 value 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 value is initially zero, and is calculated in each SPF pass.
AreaLsaCount	Specifies the total number of link state advertisements in the link state database for this area, excluding Autonomous System external link state advertisements.
AreaLsaCksumSum	Specifies the sum of the checksums of the link state advertisements contained in the link state database for this area. This sum excludes external (link state type 5) link state advertisements. The sum can determine if there has been a change in the 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. The value 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

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.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Area Aggregate** tab.
4. Click **Insert**.
5. Using the fields provided, create the new area aggregate.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Area Aggregate** tab.

Field	Description
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. Options available are: summaryLink or nssaExternalLink.

Field	Description
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 (<i>advertiseMatching</i> value) or result in the subnet not being advertised at all outside the area. Select one of the following types: <ul style="list-style-type: none"> • <i>AdvertiseMatching</i>: advertises the aggregate summary LSA with the same LSID • <i>DoNotAdvertiseMatching</i>: suppresses all networks that fall within the entire range • <i>AdvertiseDoNotAggregate</i>: advertises individual networks
AdvertiseMetric	Specifies the advertisement metric associated with this aggregate. Enter an integer value between 0 and 65535, which represents the metric cost value for the OSPF area range.

Configuring OSPF stub area metrics

Use the following procedure to view 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.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.

3. In the work area, click the **Stub Area Metrics** tab.
4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select a parameter or value from the list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Stub Area Metrics** tab.

Field	Description
Areaid	Specifies the unique ID of the stub area.
TOS	Specifies the type of service associated with the metric.
Metric	Specifies the metric value that applies 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

Use the following procedure to configure OSPF interfaces.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Interfaces** tab.
4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select a parameter or value from the list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Interfaces** tab.

Field	Description
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, or waiting.
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 field is an integer value between 0 and 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 or passive.

Field	Description
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 field is an integer value between 0 and 3600.
RetransInterval	Specifies the number of seconds between link state advertisement retransmissions for adjacencies that belong to this interface. This value is also used when retransmitting database description and link state request packets. This field is an integer value between 0 and 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 that attach to a common network. This field is an integer value between 1 and 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 that attach to the common network. This field is an integer value between 0 and 2147483647.
PollInterval	Specifies the poll interval.
AdvertiseWhenDown	Specifies whether this interface sends advertisements even when it is non-operational.
MtIgnore	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

Use the following procedure to configure OSPF interface metrics.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **If Metrics** tab.
4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select a parameter or value from the drop-down list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **If Metrics** tab.

Field	Description
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 that indicates the distance from the OSPF router to any network in the range. This field is an integer value between 0 and 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.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Message Digest** tab.
4. In the toolbar, click **Insert**.
5. Using the fields provided, create the new digest entry.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Message Digest** tab.

Field	Description
IpAddress	Specifies the IP address of the OSPF interface associated with the digest entry.
Index	Specifies an index value for the digest entry. This field is an integer value between 1 and 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.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-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 fields of the **Neighbor** tab.

Field	Description
IpAddr	Specifies the IP address this neighbor uses as an IP source address. On addressless links, this address will not appear as 0.0.0.0 but as the address of another interface that belongs to the neighbor.
AddressLessIndex	Specifies the corresponding value of the interface index on addressless links. This value is zero for interfaces that have an IP address.
RouterId	Specifies the unique ID of the neighboring router in the Autonomous System.
Options	Specifies a value that corresponds to the Options field of the neighbor.

Field	Description
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 field is a value between 0 and 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 the neighbor.

Configuring an OSPF virtual link

Use the following procedure to create an OSPF virtual link.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Virtual If** tab.
4. In the toolbar, click **Insert**.

5. Using the fields provided, create a new OSPF virtual link.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Virtual If** tab.

Field	Description
Areald	Specifies the unique ID of the area that connects 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 and 3600. The default value is 1.
RetransInterval	Specifies the number of seconds between link state advertisement retransmissions for adjacencies that belong to the virtual interface. The retransmit interval is also used for database description and link state request packets. The retransmit interval is expressed as an integer between 1 and 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 that attach to a common network. The hello interval is expressed as an integer between 1 and 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 and 2147483647. The retransmit dead interval must be a multiple of the hello interval, and must be the same for all routers that attach to a common network. The default value is 60.
AuthType	Specifies the interface authentication type. The available authentication types are: none, simplePassword, or 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.

Field	Description
Type	Specifies whether the virtual interface is broadcast or passive.

Defining MD5 keys for OSPF virtual links

Use the following procedure to configure OSPF MD5 keys for OSPF virtual interfaces.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Virtual If Message Digest** tab.
4. In the toolbar, click **Insert**.
5. Configure the parameters as required.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Virtual If Message Digest** tab.

Field	Description
Areald	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.

Field	Description
Index	Specifies the index value of the virtual interface message digest entry. The value is an integer between 1 and 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

Use the following procedure to view OSPF Virtual Neighbors information.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Virtual Neighbors** tab.
4. Click **Refresh** to update the displayed information.

Variable definitions

The following table describes the fields of the **Virtual Neighbors** tab.

Field	Description
Area	Specifies the subnetwork in which the virtual neighbor resides.
RouterId	Specifies the 32-bit integer that uniquely identifies the neighboring router in the autonomous system.
IpAddr	Specifies the IP address of the virtual neighboring router.

Field	Description
Options	Specifies a bit mask that corresponds 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.

Configuring OSPF host routes

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.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Hosts** tab.
4. IN the toolbar, click **Insert**.
 The OSPF Insert Hosts dialog box appears.
5. Configure the parameters as required.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Hosts** tab.

Field	Description
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

Use the following procedure to view OSPF link states.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Link State Database** tab.
4. In the toolbar, click **Refresh** to update the displayed information.

Variable definitions

The following table describes the fields of the **Link State Database** tab.

Field	Description
Areald	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 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	Detects 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

Use the following procedure to view the OSPF external link state database.

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.

Procedure steps

1. From the navigation tree, double-click IP Routing.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the Ext. Link State Database tab.
4. In the toolbar, click **Refresh** to update the displayed information.

Variable definitions

The following table describes the fields of the **Ext. Link State Database** tab.

Field	Description
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 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	Detects 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.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **OSPF**.
3. In the work area, click the **Stats** tab.
4. Values on the **Stats** tab will refresh automatically based on the value selected in the Poll Interval field. In the toolbar, click **Clear Counters** to clear the counters and start over at zero.

Variable definitions

The following table describes the fields of the **Stats** tab.

Field	Description
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.

Field	Description
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.

Displaying VLAN OSPF statistics

Use the following procedure to view VLAN OSPF statistical information on a per-interface basis.

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.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs**.
3. In the **Basic** tab, select an interface listed in the work area.
4. In the toolbar, click **IP**.

The IP VLAN tab appears.

5. In the work area, click the **OSPF Stats** tab.
6. In the table, select the desired data.
7. In the toolbar, click the appropriate graph button.

Variable definitions

The following table describes the fields of the **Stats** tab.

Field	Description
VersionMismatches	Specifies the number of version mismatches received by this interface.
AreaMismatches	Specifies the number of area mismatches received by this interface.
AuthTypeMismatches	Specifies the number of AuthType mismatches received by this interface.
AuthFailures	Specifies the number of authentication failures on this interface.
NetMaskMismatches	Specifies the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Specifies the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Specifies the number of dead interval mismatches received by this interface.
OptionMismatches	Specifies the number of option mismatches received by this interface.
RxHellos	Specifies the number of hello packets received by this interface.
RxDBDescrs	Specifies the number of database descriptor packets received by this interface.
RxLSUpdates	Specifies the number of link state update packets received by this interface.
RxLSReqs	Specifies the number of link state request packets received by this interface.
RxLSAcks	Specifies the number of link state acknowledge packets received by this interface.
TxHellos	Specifies the number hello packets transmitted by this interface.
TxDBDescrs	Specifies the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Specifies the number of link state update packets transmitted by this interface.
TxLSReqs	Specifies the number of link state request packets transmitted by this interface.
TxLSAcks	Specifies the number of link state acknowledge packets transmitted by this interface.

Chapter 17: 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 brouter port that you want to enable with RIP.
Routing is automatically enabled on the VLAN when you assign an IP address to it.

RIP configuration procedures

To configure RIP routing on the Ethernet Routing Switch, perform the following steps:

Procedure steps

1. Enable RIP globally.
2. Configure global RIP properties as required.
3. Enable RIP on the desired VLAN or brouter interfaces.
4. Configure interface RIP properties as required.

Configuring the global RIP status

Use this procedure to globally enable RIP on the switch.

Procedure steps

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter 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 the RIP global timeout, holddown timer, and update timer

Use this procedure to set the RIP global timeout, holddown timer, and update timer.

Procedure steps

1. Log on to the RIP Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] timers basic [holddown <holddown-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:

Variable	Value
	<ul style="list-style-type: none"> • holddown timer: 120 seconds • global timeout: 180 seconds • update timer: 30 seconds
<holddown-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 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 Ethernet Routing 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.

Procedure steps

1. Log on to the RIP Router Configuration mode in ACLI.
2. At the command prompt, enter 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.
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.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter 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 the RIP status on an interface

Use this procedure to configure the RIP status on a VLAN interface or router port.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] ip rip enable
```

OR

3. Log on to the Router Configuration mode in ACLI.
4. At the command prompt, enter the following command:

```
[no] network <ip_address>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Disables RIP on the interface.
[no]	Disables RIP on the IP interface.
<ip_address>	The IP address of the interface to be configured.

Configuring RIP parameters for an interface

Use this procedure to configure RIP parameters for an interface.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] ip rip [advertise-when-down enable] [auto-
aggregation enable] [cost <cost>] [default-listen enable]
[default-supply enable] [holddown {<holddown> | global}]
[listen enable] [poison enable] [proxy-announce enable]
[receive version {rip1 | riplorrip2 | rip 2}] [send version
```

```
{rip1 | rip1orrip2 | 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 the specified configuration from the switch.
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 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.
holddown	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. If the interface timer is configured, this setting overrides the global parameter and does not change if the global parameter is modified. Range is 0–360 seconds. The default value is set by the global holddown parameter, which has a default of 120 seconds.
listen enable	Enables this interface to listen for RIP advertisements. The default value is enabled.

Variable	Value
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 {rip1 rip1orrip2 rip 2}	Specifies the RIP version received on this interface. Default is rip1orrip2.
send {notsend rip1 rip1comp rip 2}	Specifies the RIP version sent on an interface. Default is rip1compatible.
supply enable	Enables RIP router advertisements on this interface. The default value is enabled.
timeout <timeout>	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. Range is 15–259200 seconds. The default value is set by the global timeout parameter, which has a default of 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 RIP interface configuration

Use this procedure to display configuration for a RIP interface.

Procedure steps

1. Log on to the XXXX Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip rip [interface [<vid>] [FastEthernet [<portlist>]]
[VLAN [<vid>]] ]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[interface]	Displays RIP statistics by interface. Omission of this key word displays general RIP information
[<vid>]	Displays RIP information for the specified VLAN.
[FastEthernet [<portlist>]]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[VLAN [<vid>]]	Displays RIP information for the specified VLAN. 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.
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.

Field	Description
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.

Triggering a RIP update manually

Use this procedure to manually trigger a RIP update on an interface.

Procedure steps

1. Log on to the XXXX Configuration mode in ACLI.
2. At the command prompt, enter the following command:
`manualtrigger ip rip interface vlan <vid>`

Chapter 18: RIP configuration examples using ACLI

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

RIP is configured 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. This practice does not imply that a brouter port is the only means through which a core connection can be established. The use of a brouter port is only one of many ways to create such a connection.

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.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 51 name "VLAN-51" type port
5530-24TFD(config)# interface vlan 51
5530-24TFD(config-if)# ip address 10.10.1.1 255.255.255.0
5530-24TFD(config-if)# exit
5530-24TFD(config)# vlan members add 51 8-9
```

2. Enable RIP using one of the following command sequences.

```
5530-24TFD(config)# interface vlan 51
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# exit
```

OR

```
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 10.10.1.1
5530-24TFD(config-router)# exit
```

3. Select the VLAN to configure RIP interface properties.

```
5530-24TFD(config)# interface vlan 51
```

4. Disable Supply RIP Updates on the VLAN, if required.

```
5530-24TFD(config-if)# ip rip supply disable
```

5. Disable Listen for RIP Updates on the VLAN, if required.

```
5530-24TFD(config-if)# ip rip listen disable
```
6. Enable Default Route Supply on the VLAN, if a default route exists in the route table.

```
5530-24TFD(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.

```
5530-24TFD(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).

```
5530-24TFD(config-if)# ip rip out-policy map1
```
9. Enable Triggered Updates on the VLAN, if required.

```
5530-24TFD(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.

```
5530-24TFD(config-if)# ip rip cost 2
```
11. Configure send mode parameters on the VLAN.

```
5530-24TFD(config-if)# ip rip send version rip2
```
12. Configure receive mode parameters on the VLAN.

```
5530-24TFD(config-if)# ip rip receive version rip2
```
13. Enable poison reverse on the VLAN.

```
5530-24TFD(config-if)# ip rip poison enable
```

Configuring RIP

This section describes a basic RIP configuration setup between two Avaya Ethernet Routing Switch 5000 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.

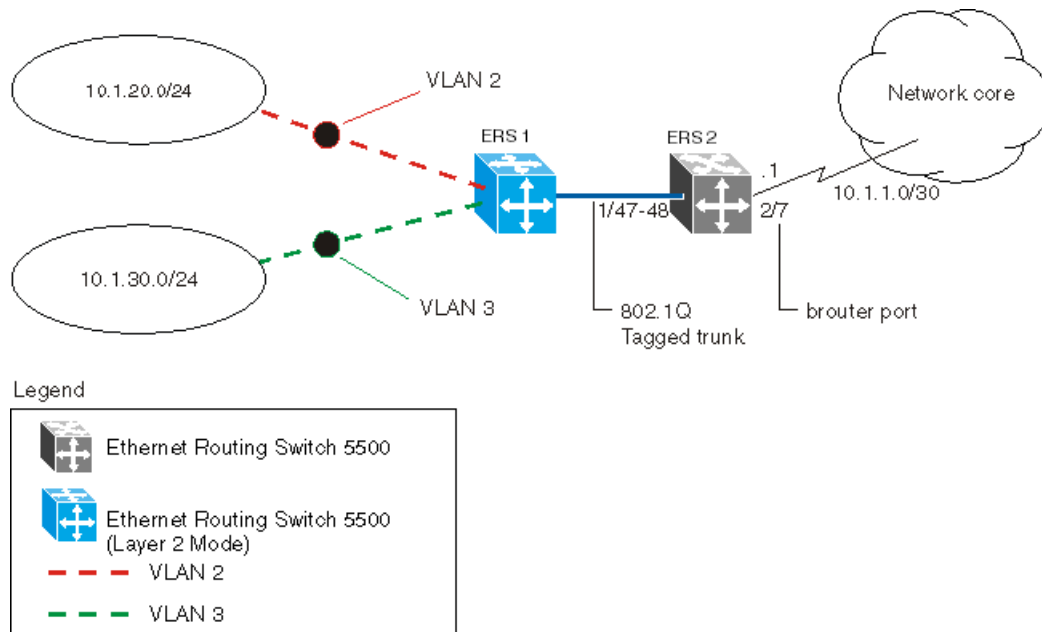


Figure 42: RIP configuration example

For the purposes of 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 broturer port with RIP to connect to the network core.

Use the following procedure to configure router ERS 2 and reproduce the illustrated RIP configuration:

1. Configure tagging on ports 1/47 and 1/48. Tagging is required to support multiple VLANs on the same interface.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# vlan ports 1/47-48 tagging tagAll
```

2. Configure ERS2 for VLAN 2 access.

- a. Create a port-based VLAN (VLAN 2) using spanning tree group 1 and include ports 1/47 and 1/48 in VLAN 2.

```
5530-24TFD(config)# vlan create 2 name "VLAN-2" type
port
5530-24TFD(config)# vlan member add 2 port 1/47-48
```

- b. Assign the IP address 10.1.20.2/24 to VLAN 2.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.2
255.255.255.0
```

- c. 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.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip supply disable
5530-24TFD(config-if)# ip rip listen disable
```

3. Configure ERS2 for VLAN 3 access.

- a. Create a port-based VLAN (VLAN 3) using spanning tree group 1 and include ports 1/47 and 1/48 in VLAN3.

```
5530-24TFD(config)# vlan create 3 name "VLAN-3" type
port
5530-24TFD(config)# vlan member add 3 port 1/47-48
```

- b. Assign the IP address 10.1.30.2/24 to VLAN 3.

```
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip address 10.1.30.2 255.255.255.0
```

- c. 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.

```
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip supply disable
5530-24TFD(config-if)# ip rip listen disable
```

4. 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.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# brouter vlan 2090 subnet 10.1.1.1/30
```

* Note:

Usage of the **brouter** command above requires the usage of Variable Length Subnetting. Usage of a dotted decimal subnet mask is not allowed.

- b. Enable RIP on the interface.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# ip rip enable
```

5. Enable IP routing and RIP globally.

```
5530-24TFD(config)# ip routing
5530-24TFD(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 fastEthernet 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 5000 Series can be configured to operate using RIP version 1 or 2. The configuration illustrated below demonstrates an Avaya Ethernet Routing Switch 5000 Series switch that has been configured to operate use RIP version 2 only.

*** Note:**

This example builds on the previous RIP configuration.

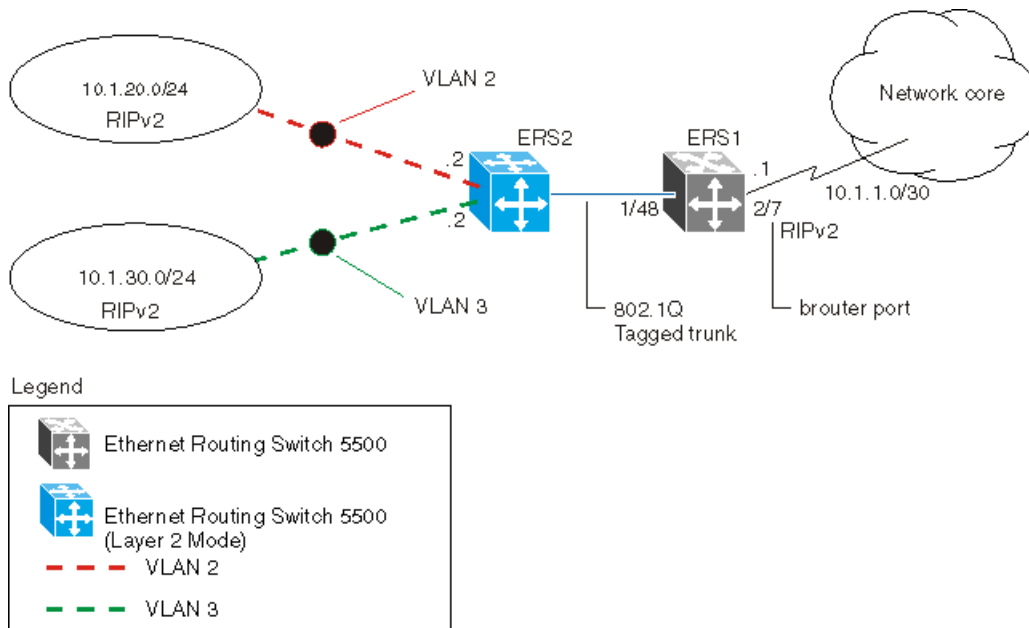


Figure 43: RIPv2 configuration example

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.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip rip send version rip2
5530-24TFD(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.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip rip send version rip2
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 3
5530-24TFD(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.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip send version rip2
5530-24TFD(config-if)# ip rip receive version rip2
```

Using RIP accept policies

RIP accept policies are used on the Avaya Ethernet Routing Switch 5000 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 5000 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

RIP configuration examples using ACLI

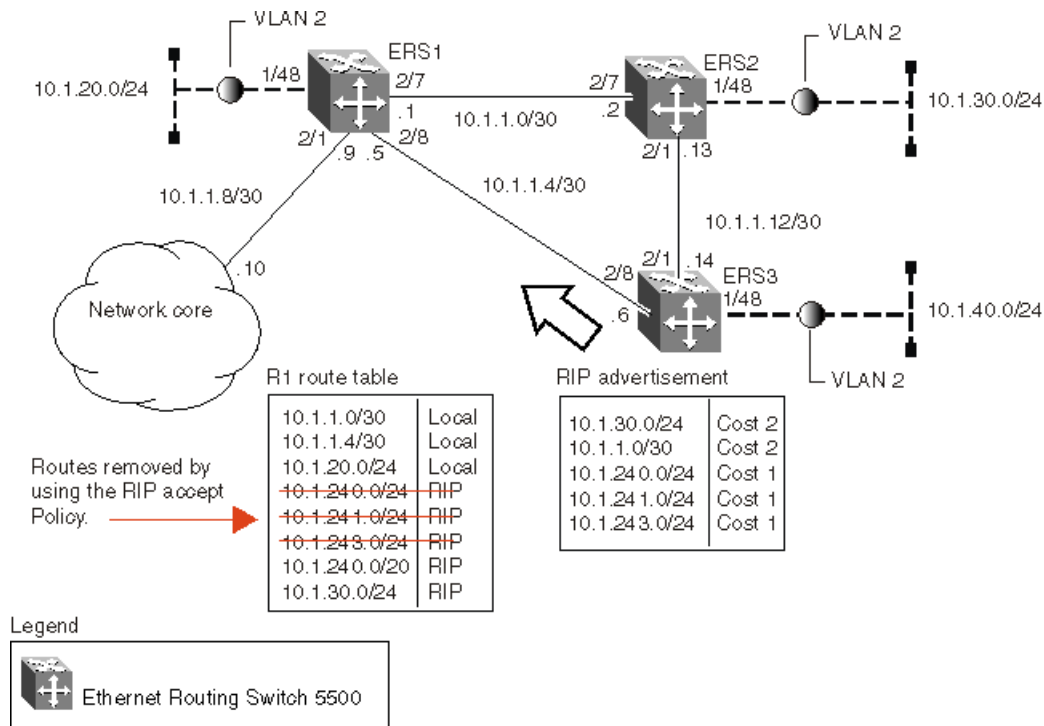


Figure 44: Accept policy configuration

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.

```
5530-24TFD(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.

```
5530-24TFD(config)# route-map rip_pol_1 1
5530-24TFD(config)# route-map rip_pol_1 1 enable
5530-24TFD(config)# route-map rip_pol_1 permit 1 enable
5530-24TFD(config)# route-map rip_pol_1 permit 1 match network Prefix_1
5530-24TFD(config)# route-map rip_pol_1 permit 1 set-injectlist Prefix_1
```

3. Add the route policy created in step 2 to both RIP core ports.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# brouter vlan 2090 subnet 10.1.1.1/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip in-policy rip_pol_1
5530-24TFD(config)# interface FastEthernet 2/8
5530-24TFD(config-if)# brouter vlan 2091 subnet 10.1.1.5/30
5530-24TFD(config-if)# ip rip enable 5530-24TFD(config-if)# ip rip in-policy rip_pol_1
```

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:

```

rip_pol_1
! *** Route Policies *** !
ip prefix-list Prefix_1 10.1.240.0/20 ge 20 le 32
route-map rip_pol_1
route-map rip_pol_1 1 enable
no route-map rip_pol_1 1 match interface
route-map rip_pol_1 1 match metric 0
route-map rip_pol_1 1 match network Prefix_1
no route-map rip_pol_1 1 match next-hop
route-map rip_pol_1 1 match route-type any
no route-map rip_pol_1 match route-source
route-map rip_pol_1 1 set injectlist Prefix_1
route-map rip_pol_1 set mask 0.0.0.0
route-map rip_pol_1 set metric 0
route-map rip_pol_1 set nssa-pbit enable
route-map rip_pol_1 set ip-preference 0
! *** Brouter Port *** !
interface fastEthernet ALL
brouter port 2/7
vlan 2090 subnet 10.1.1.1/30
ip rip in-policy rip_pol_1
brouter port 2/8 vlan 2091 subnet 10.1.1.5/30
ip rip in-policy rip_pol_1

```

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.

```

5530-24TFD(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.

```

5530-24TFD(config)# route-map rip_pol_1 1
5530-24TFD(config)# route-map rip_pol_1 1 enable
5530-24TFD(config)# route-map rip_pol_1 permit 1 enable
5530-24TFD(config)# route-map rip_pol_1 permit 1 set-injectlist Prefix_1

```

3. Add the route policy created in step 2 to both RIP core ports.

```

5530-24TFD(config)# interface FastEthernet 2/1
5530-24TFD(config-if)# brouter vlan 2091 subnet 10.1.1.14/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip out-policy rip_pol_1
5530-24TFD(config)# interface FastEthernet 2/8
5530-24TFD(config-if)# brouter vlan 2090 subnet 10.1.1.6/30

```

```
5530-24TFD(config-if)# ip rip enable
5530-24TFD(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.

```
5530-24TFD(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.

```
5530-24TFD(config)# route-map rip_pol_2 deny 1 enable match network
Prefix_2
5530-24TFD(config)# route-map rip_pol_2 1 match network Prefix_2
```

3. Add the Route Policy created in step 2 to both RIP core ports.

```
5530-24TFD(config)# interface FastEthernet 2/1
5530-24TFD(config-if)# brouter vlan 2091 subnet 10.1.1.14/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip out-policy rip_pol_2
5530-24TFD(config)# interface FastEthernet 2/8
5530-24TFD(config-if)# brouter vlan 2090 subnet 10.1.1.6/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip out-policy rip_pol_2
```


Chapter 19: RIP configuration using Enterprise Device Manager

This chapter describes the procedure you can use to configure and manage the Routing Information Protocol (RIP) on the Avaya Ethernet Routing Switch 5000 Series . 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 would be 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 RIP using EDM

Use the following procedure to configure RIP routing on the Ethernet Routing Switch.

Procedure steps

1. Enable RIP globally.
2. Configure global RIP properties as required.
3. Enable RIP on the desired VLAN or brouter interfaces.
4. Configure interface RIP properties as required.

Configuring Global RIP properties using EDM

Use the following procedure to configure global RIP parameters.

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.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **RIP**.
3. In the RIP work area, click the **Globals** tab.
4. Configure global RIP parameters as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** to verify the global RIP configuration.

Variable definitions

The following table describes the fields of the **Globals** tab.

Field	Definition
Operation	Enables or disables the operation of RIP on all interfaces. The default is disabled.
UpdateTime	Indicates the time interval between RIP updates on all interfaces. It is a global parameter for the box; that is, it applies to all interfaces and cannot be set individually for each interface. The default is 30 seconds.
RouteChanges	Indicates the number of route changes made to the IP Route Database by RIP; does not include the refresh of a route's age.
Queries	Indicates 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 to 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

Field	Definition
	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.

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.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **RIP**.
3. In the RIP work area, click the **Interface** tab.
4. Configure the RIP interface as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** to verify the RIP interface configuration.

Variable definitions

The following table describes the fields of the **Interface** tab.

Field	Definition
Address	Indicates the IP address of the RIP interface. This field is for organizational purposes only and cannot be edited.
Send	<p>Sets the RIP version sent on this interface. The following values are valid:</p> <ul style="list-style-type: none"> • 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. <p>The default is rip1Compatible.</p>
Receive	<p>Sets the RIP version received on this interface: rip1, rip2, or rip1OrRip2. The default is rip1OrRip2. Note that rip2 and rip1OrRip2 imply reception of multicast packets.</p>

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.

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.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **RIP**.
3. In the RIP work area, click the **Interface Advance** tab.
4. Configure the advanced RIP parameters as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** to verify the advanced RIP configuration.

Variable definitions

The following table describes the fields of the **Interface Advance** tab.

Field	Definition
Address	Indicates the IP address of the RIP interface. This field is for organizational purposes only and cannot be edited.
Interface	Indicates 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 router advertisements.
TriggeredUpdate	Enables or disables triggered updates on this interface.
AutoAggregate	Enables or disables auto aggregation on this interface.
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	Indicates the cost associated with this interface.
HoldDownTime	Sets the holddown timer for this interface. This is an integer value in seconds between 0 and 360.
TimeoutInterval	Sets the timeout interval for this interface. This is an integer value between 15 and 259200.
ProxyAnnounceFlag	Enables or disables proxy announcements on this interface.

Displaying RIP Statistics using EDM

Use the following procedure to view RIP statistics.

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.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, double-click **RIP**.
3. In the RIP work area, click the **Stats** tab.
4. To select a record to graph, click a table row.
5. On the toolbar, click **Graph**.
6. On the toolbar, click the **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart** icon, to select a graph type.

Variable definitions

The following table describes the fields of the **RIP—Stats** tab.

Field	Definition
Address	Indicates the RIP interface address.
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.

Configuring RIP parameters for a VLAN using EDM

Use the following procedure to configure VLAN RIP parameters.

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.

Procedure steps

1. From the navigation pane, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **Basic** tab.
4. To select a VLAN, click a table row.
5. On the toolbar, click **IP**.
6. In the IP, VLAN work area, click the **RIP** tab.
7. Configure RIP parameters as required.
8. On the toolbar, click **Apply**.
9. On the toolbar, you can click **Refresh** to verify the VLAN RIP configuration.

Variable definitions

The following table describes the fields of the **RIP** tab.

Field	Definition
Poison	Determines whether or not poison reverse is implemented on this interface.
DefaultSupply	Determines whether or not the interface implements the default supply mechanism.
DefaultListen	Determines whether or not the interface implements the default listen mechanism.
AutoAggregateEnable	Determines whether or not auto aggregation is enabled on this interface.
AdvertiseWhenDown	Determines whether or not this interface will advertise even when non-operational.

Field	Definition
Cost	Indicates the cost associated with this interface. Value ranges between 1 and 15. Default value is 1.

Chapter 20: ECMP configuration using ACLI

This chapter describes the procedures you can use to configure Equal Cost MultiPath (ECMP) using Avaya Command Line Interface (ACLI).

The ECMP feature allows routers to determine equal cost paths to the same destination prefix. The multiple paths can be used for load sharing of traffic and allows faster convergence to other active paths in case of network failure.

Important:

Although ECMP can function in a mixed stack environment, the Avaya ERS 5510 does not support ECMP, so the ECMP feature cannot run on any Avaya ERS 5510 units in a mixed stack.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Configure routing (RIP, OSPF, BGP, or static routes) on the switch.

Configuring the number of ECMP paths allotted for RIP

Use this procedure to configure the number of ECMP paths allotted for the Routing Information Protocol (RIP).

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] rip maximum-path <path_count>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets the maximum ECMP paths allowed to the default value, 1.
[no]	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Configuring the number of ECMP paths allotted for OSPF

Use this procedure to configure the number of ECMP paths allotted for the Open Shortest Path First (OSPF) protocol.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] ospf maximum-path <path_count>
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Sets the maximum ECMP paths allowed to the default value, 1.
[no]	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Configuring the number of ECMP paths allotted for static routes

Use this procedure to configure the number of ECMP paths allotted to static routes.

Procedure steps

To configure the number of ECMP paths allotted to static routes enter the following from the Global Configuration mode:

```
[default] [no] maximum-path <path_count>
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Sets the maximum ECMP paths allowed to the default value, 1.
no	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Configuring the number of ECMP paths allotted for BGP

Use this procedure to configure the number of ECMP paths allotted to Border Gateway Protocol (BGP).

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] [no] bgp maximum-path <path_count>
```

Variable definitions

The following table describes the parameters for the `bgp maximum-path` command.

Variable	Value
default	Sets the maximum ECMP paths allowed to the default value, 1.
no	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Displaying ECMP path information

Use this procedure to display ECMP path information.

Procedure steps

1. Log on to the User Exec mode in ACLI.
2. At the command prompt, enter the following command:

```
show ecmp
```

Job aid

The following table shows the field descriptions for the `show ecmp` command.

Field	Description
Protocol	Indicates the protocol.
MAX-PATH	Indicates the maximum number of equal-cost paths supported for the listed protocol.

ECMP configuration examples

Equal Cost Multipath (ECMP) is an IP feature for load-balancing routed IP traffic across up to four equal-cost paths for each supported protocol. ECMP supports OSPF, RIP, and static routes. Some benefits of using ECMP:

- Supported protocols will rerun when an ECMP path fails, and the other configured paths will automatically take the load.
- Load sharing implies better use of network facilities.

ECMP is selected based on the source and destination IP address in the packet. The hash_control register has a HASH_SELECT field which is set to 5 (lower CRC-32).

$R1 = \text{CRC32}(\text{SIP}, \text{DIP})$

$R2 = R1 \& 0x1F$ (The Least Significant 5 bits are selected)

$\text{ecmp_index} = R2 \% (\text{ecmp_count} + 1)$

*** Note:**

The value `ecmp_count` above is zero-based in the hardware so if four paths are present then the value is three. This is why the value is `ecmp_count + 1`.

The ECMP traffic distribution algorithm is demonstrated in the following example:

Consider two network devices, Device 1 at the IP address 192.1.1.3 and Device 2 at 192.1.1.4. Device 1 send to Device 2 so that 192.1.1.3 is the source IP address (SIP) and 192.1.1.4 is the destination IP address (Device 2).

To calculate the CRC32 for the example source and destination IP address noted above, the following calculations would be made:

- CRC32 polynomial : $x^{32} + x^{28} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x^1 + 1$
- $R1 = \text{CRC32}(0xc0010103, 0xc0010104) = 0xf474b549$
- $R2 = (0xf474b549 \& 0x1f) = 9$

If, for the purposes of this example, it is assumed that the ECMP count is 4 (hardware entries 0 though 3), the following calculation is then made:

- $\text{ecmp_index} = 9 \% (4+1) = 1$

This means that in this example, the second path at hardware index 1 in the ECMP table will be used.

In the configuration example below, the following command would enable two OSPF ECMP paths on router R1:

```
5530-24TFD(config)#ospf maximum-path 2
```




Figure 45: ECMP configuration example

Use the following commands to enable ECMP on each of the supported protocols:

- **OSPF**

```
ospf maximum-path <path_count>
```
- **RIP**

```
rip maximum-path <path_count>
```
- **Static Routes**

```
maximum-path <path_count>
```
- **BGP**

```
bgp maximum-path <path_count>
```

In all commands above, the *<path_count>* parameter represents the number of ECMP paths allotted. Value range is 1–4 and the default value is 1.

Displaying the IP routing table

After ECMP configuration is complete, verify the ECMP paths in the routing table using the `show ip route` command. The following example displays the output for this command:

```

=====
Ip Route
=====
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
  
```

```

-----
0.0.0.0 0.0.0.0 10.100.111.110 119 SIB 5
3.3.3.0 255.255.255.0 3.3.3.112 - C DB 0
4.4.4.0 255.255.255.0 4.4.4.112 - C DB 0
5.5.5.0 255.255.255.0 5.5.5.1 1 2 - C DB 0
10.10.10.0 255.255.255.0 10.10.10.1 1 5 - C DB 0
10.100.111.0 255.255.255.0 10.100.111.200 11- C DB 0
Total Routes: 6
-----
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. In this example, two routes to IP address 10.1.40.0 and two routes to IP address 10.1.30.0 are displayed.

Displaying global ECMP configuration

To confirm global ECMP configuration, use the `show ecmp [vrf <vrf-name>] [vrfs <vrf-ids>]` command. A sample output from this command is displayed below:

```

5650TD-PWR# show ecmp
MAX-PATH
-----
static:          1
rip:             2
ospf:           4
bgp:            1

```


Chapter 21: ECMP configuration using Enterprise Device Manager

This chapter describes the procedure you can use to configure Equal Cost MultiPath (ECMP) using Enterprise Device Manager (EDM).

The ECMP feature allows routers to determine equal cost paths to the same destination prefix. The multiple paths can be used for load sharing of traffic and allows faster convergence to other active paths in case of network failure.

Important:

Although ECMP can function in a mixed stack environment, the Avaya ERS 5510 does not support ECMP, so the ECMP feature cannot run on any Avaya ERS 5510 units in a mixed stack.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Configure routing (RIP, OSPF, BGP, or static routes) on the switch.

Configuring ECMP using EDM

Use the following procedure to configure and 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 pane, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the work area, click the **ECMP** tab.
4. To select a protocol for which to configure ECMP, double-click the table cell in the **MaxPath** column heading for the protocol.
5. Type a value in the text box.
6. Repeat steps **4** and **5** as required.
7. On the toolbar, click **Apply**.
8. On the toolbar, you can click **Refresh** to verify the ECMP configuration.

Variable definitions

The following table describes the fields of the ECMP tab.

Field	Description
RoutingProtocol	Indicates the routing protocol to be configured.
MaxPath	Indicates the maximum number of ECMP paths assigned to the protocol. Value ranges between 1 and 4. Default is 1.

Chapter 22: Route policy 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 procedures

To configure routing policies, perform the following steps:

1. Create the appropriate prefix lists.
2. Assign those prefix lists to route maps.
3. Apply the route maps to the appropriate type of policy.

Configuring prefix lists

Use this procedure to configure up to four prefix lists for use in route policies.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter 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.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] route-map <map_name> [permit|deny] <sequence_number>
[enable] [match {as-path <WORD> | community <WORD> |
community-exact <enable> | interface <prefix_list> | local-
preference <lp_value> | 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 {as-path <WORD> |
as-path-mode <prepend | tag> | community <WORD> | community-
mode <additive | none | unchanged> | injectlist <prefix_list>
| ip-preference <pref> | local-preference <0-65535> | mask
<ip_address> | metric <metric_value> | metric-type
```

```
<metric_type> next-hop <A.B.C.D> | nssa-pbit enable | origin
<egp | igp | incomplete> | weight <0-65535>}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
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.
enable	Specifies whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored.
[match {as-path <WORD> community <WORD> community-exact <enable> interface <prefix_list> local-preference <lp_value> metric <metric_value> network <prefix_list> next-hop <prefix_list> protocol <protocol_name> route-source <prefix_list> route-type <route_type>}]	<p>If configured, the switch matches the specified criterion:</p> <ul style="list-style-type: none"> • as-path <WORD>: matches the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified AS-lists. This field is used only for BGP routes and ignored for all other route types. <WORD> specifies the list IDs of up to four AS—lists, separated by a comma. • community <WORD>: matches the community attribute of the BGP routes against the contents of the specified community lists. This field is used only for BGP routes and ignored for all other route types. <WORD> specifies the list IDs of up to four defined community lists, separated by a comma. • community-exact <enable>: <ul style="list-style-type: none"> - when enabled, match-community-exact results in a match when the community attribute of the BGP routes matches all of the entries of all the community lists specified in match-community. - when disabled, match community-exact results in a match when the community attribute of the BGP routes match any entry of any community-list specified in match-community. • interface <prefix_list>: matches the IP address of the received interface against the contents of the specified prefix list.

Variable	Value
	<ul style="list-style-type: none"> • local-preference <0–2147483647>: matches the local preference, applicable to all protocols. • metric <metric_value>: matches the metric of the incoming advertisement or existing route against the specified value, an 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 {as-path <WORD> as-path-mode {prepend tag} community <WORD> community-mode <additive none unchanged> injectlist <prefix_list> ip-preference <pref> local-preference <0–65535> mask <A.B.C.D> metric <metric_value> metric-type <metric_type> next-hop <A.B.C.D> nssa-pbit enable origin <egp igp incomplete> weight <0–65535>}]	<p>If configured, the switch sets the specified parameter:</p> <ul style="list-style-type: none"> • as-path <WORD>: adds the AS number of the AS-list to the BGP routes that match this policy. <WORD> specifies the list ID of up to four defined AS-lists, separated by a comma. <p>! Important: Configuring the set-as-path attribute does not change the AS path information for IBGP peers.</p> <ul style="list-style-type: none"> • community <WORD>: when configured, the switch adds the community number of the community list to the BGP routes that match this policy. • community-mode <additive none unchanged>: configures the community mode. <ul style="list-style-type: none"> - additive — the switch prepends the community number of the community list specified in set-community to the old community path attribute of the BGP routes that match this policy. - none — the switch removes the community path attribute of the BGP routes that match this policy to the specified value.

Variable	Value
	<ul style="list-style-type: none"> - unchanged — the switch retains the previously configured BGP route attribute. • 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. • local-preference <0–65535>: a value used during the route decision process in the BGP protocol. Applicable to BGP only. • mask <A.B.C.D>: sets the mask 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. • next-hop <A.B.C.D>: specifies the IP address of the next-hop router. • nssa-pbit enable: enables the NSSA N/P-bit, which notifies the ABR to export the matching external route. Used for OSPF policies only. • origin <egp igp incomplete>: when configured, the switch changes the origin path attribute of the BGP routes that match this policy to the specified value. <p>! Important:</p> <p style="padding-left: 20px;">The set-origin attribute applies to only outbound route policies.</p> <ul style="list-style-type: none"> • weight <0–65535>: specifies the weight value for the routing table. Used for BGP only. A value of 0 indicates that this parameter is not set.

Displaying route maps

Use this procedure to display configured route maps.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter 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.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. To specify a RIP Accept policy for an interface, enter the following command at the command prompt:

```
[default] [no] ip rip in-policy <rmap_name>
```

3. Log on to the User EXEC mode in ACLI.
4. To display RIP interface configuration, enter the following command at the command prompt:

```
show ip rip [interface [<vid>] [FastEthernet [<portlist>]]  
[VLAN [<vid>]] ]
```

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.
[interface]	Displays RIP statistics by interface. Omission of this key word displays general RIP information
[<vid>]	Displays RIP information for the specified VLAN.
[FastEthernet [<portlist>]]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[VLAN [<vid>]]	Displays RIP information for the specified VLAN. If no VLAN ID is specified, all VLAN information is displayed.

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.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. To apply a RIP Announce (Out) policy to an interface, enter the following command at the command prompt:

```
[default] [no] ip rip out-policy <rmap_name>
```

3. Log on to the User EXEC mode in ACLI.
4. To display RIP interface configuration, enter the following command at the command prompt:

```
show ip rip [interface [<vid>] [FastEthernet [<portlist>]]
[VLAN [<vid>]] ]
```

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.
[interface]	Displays RIP statistics by interface. Omission of this key word displays general RIP information
[<vid>]	Displays RIP information for the specified VLAN.
[FastEthernet [<portlist>]]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[VLAN [<vid>]]	Displays RIP information for the specified VLAN. If no VLAN ID is specified, all VLAN information is displayed.

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.

Procedure steps

1. Log on to the OSPF Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[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
[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 <i>0.0.0.0</i> denotes that advertisements from all routers are accepted.
enable	Enables the accept entry for the router specified in the <i><ip_address></i> 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.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip ospf apply accept
```

Displaying the OSPF accept policy

Use this procedure to display the OSPF accept policy.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip ospf accept
```

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.

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.

Procedure steps

1. Log on to the Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
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
[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.
<rmap_name>	Specifies the route policy to associate with route redistribution. This is the name of a previously configured route map.

Variable	Value
<metric_value>	Specifies the metric value to associate with the route redistribution. This is an integer value between 0 and 65535.
<metric_type>	Specifies the metric type to associate with the route redistribution. Valid options are type1 and type2.
<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.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter 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 configuration to the switch
rip	Applies only RIP OSPF redistribution configuration on the switch.
static	Applies only static OSPF redistribution configuration on the switch.

Displaying the OSPF redistribution policy

Use this procedure to display the OSPF redistribution policy configuration and status.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip ospf redistribute
```

Configuring IP forwarding next-hop

Use this procedure to configure a policy for IP forwarding next-hop.

Prerequisites

- You must enter ACLI commands through the Base Unit.
- You must always select the longest subnet mask. Network mask cannot override the longest match. For example, subnets 10.0.0.0/8 and 10.10.0.0/16 cannot apply for the same VLAN.
- You can create up to a maximum of 4 IP Forwarding Next-hop policies.
- Depending on hardware filter resource availability, up to 16 IP Forwarding Next-hop instances are allowed. Other applications, such as QoS, can consume hardware filters.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. To globally enable the IP forwarding next-hop feature, enter the following from Configuration mode:

```
ip fwd-nh [enable]
```
3. To disable the IP forwarding next-hop feature, enter the following from Configuration mode:

```
no ip fwd-nh
```
4. To create an IP forwarding next-hop policy, enter the following command at the command prompt:

```
ip fwd-nh policy <policy-name> match <source-ip/mask> [port-
type <both|tcp|udp>] [port-min <0-65535> port-max <0-65535>]
set next-hop <next-hop> [secondary-next-hop <sec-next-hop>]
```

5. Log on to the Interface Configuration mode in ACLI.
6. To apply an IP forwarding next-hop policy to a VLAN, enter the following command at the command prompt:

```
ip fwd-nh policy <policy-name> [mode {drop | normal-routing}]
```

7. To enable or disable IP forwarding next-hop policy data on a VLAN, enter the following command at the command prompt:

```
ip fwd-nh admin-status {enable|disable}
```

8. Log on to the Global Configuration mode in ACLI.
9. To delete a policy, enter the following command at the command prompt:

```
no ip fwd-nh policy <policy-name>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<policy-name>	Name of the next-hop forwarding policy. The value consists of any alphanumeric values, ranging from 1 to 32 characters.
<source-ip/mask>	Source IP address and mask length to be matched. IP address 0.0.0.0 is not supported.
<next-hop>	Next hop IP address to be used to forward the packet. The next hop must be a direct connection to any of the routing interfaces of the switch.
<sec-next-hop>	Secondary next hop IP address to be used to forward the packet.
[mode {drop normal-routing}]	Specifies the packet forwarding decision to be made based when the next-hop is not reachable. <ul style="list-style-type: none"> • drop: if the next-hop is not reachable, packets are dropped. • normal-routing: if the next-hop is not reachable, the packet follows the normal routing. This is the default value.

Displaying the IP forwarding next-hop configuration

Use this procedure to display the IP forwarding next-hop configuration.

Procedure steps

1. Log on to the Privileged EXEC mode in ACLI.
2. To display the global status of the IP forwarding next-hop feature, enter the following command at the command prompt:

```
show ip fwd-nh
```

3. To display the IP forwarding next-hop policy configuration, enter the following command at the command prompt:

```
show ip fwd-nh policy {<policy-name> | interface [vlan  
<vid>]}
```

Variable definitions

The following table describes the command variables.

Variable	Value
<policy-name>	Name of the next-hop forwarding policy. The value consists of any alphanumeric values.
<vlan-id>	Specifies a VLAN ID for which to display next-hop forwarding policies.

Chapter 23: Route policy 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.

Configuring route policies

Use the following procedure to configure routing policies.

1. Create the appropriate prefix lists.
2. Assign those prefix lists to route policies.
3. Apply the route policies to the appropriate policy type.

Creating a prefix list

Use the following procedure to create a new prefix list.

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.

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. Using the fields provided, create a new prefix list.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Prefix List** tab.

Field	Definition
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 higher 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. Default value is the mask length (PrefixMaskLen).

Creating a route policy

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. In the toolbar, click **Insert**.
5. Using the fields provided, create a new prefix list.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Route Policy** tab.

Field	Definition
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. A value of permit indicates that the route is allowed while 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. Options are RIP, Static, Direct, OSPF or 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). Available options are: <ul style="list-style-type: none"> • any • local • internal • external • externaltype1 • externaltype2

Field	Definition
	Externatype1and Externatype2 specify the OSPF routes of the specified type only.
MatchMetric	If configured, matches the metric of the incoming advertisement or existing route against the specified value (1 to 655535). If 0, then this field is ignored. The default is 0.
MatchAsPath	Configures if the system matches the BGP autonomous system path. Applicable to BGP only. This overrides the BGP neighbor filter list information.
MatchCommunityExact	Indicates if the match must be exact (that is, all of the communities specified in the path must match). Applicable to BGP only. The default is disabled.
MatchCommunity	Filters incoming and outgoing updates based on a community list. Applicable to BGP only.
MatchLocalPref	Configures if the system matches the local preference. Applicable to BGP only. 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. You can set a value from 0 to 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.
SetNextHop	Configures the IP address of the next-hop router. Applicable to BGP only.
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.

Field	Definition
SetAsPath	Indicates the AS path value to use whether the SetAsPathMode field is Tag or Prepend. Applicable to BGP only.
SetASPathMode	Configures if the system converts the tag of a route into an AS path. Applicable to BGP protocol only. The mode is either Tag or Prepend tag. The value is applicable only while redistributing routes to BGP.
SetCommunityNumber	Configures the community number for BGP advertisements. This value can be a number (1 to 42949672000) or no-export or no-advertise.
SetCommunityMode	<p>Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged.</p> <ul style="list-style-type: none"> • unchanged — keeps the community attribute in the route path as it is • none — removes the community in the route path additive • append — adds the community number specified in SetCommunityNumber to the community list attribute
SetOrigin	Configures the origin for the BGP protocol to IGP, EGP, incomplete, or unchanged. If not configured, the system uses the route origin from the IP routing table (protocol). The default is unchanged.
SetLocalPref	Configures the local preference for the BGP protocol only. The system uses this value during the route decision process for the BGP protocol. The default is 0.
SetWeight	Applicable to BGP protocol only. This field must be used with the match as-path condition. It is the weight value for the routing table. For BGP, this value overrides the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. The default is 0.

Configuring RIP in and out policies

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 **Route In/Out Policy** tab.
4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select a parameter or value from the drop-down list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **RIP In/Out Policy** tab.

Field	Definition
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

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. In the toolbar, click **Insert**.
5. Use the fields provided to create the new accept policy.
6. Click **Insert**.
7. Using the fields provided, configure the desired accept policy.
8. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **OSPF Accept** tab.

Field	Definition
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.
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

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. In the toolbar, click **Insert**.
5. Using the fields provided, create the new redistribution entry.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Redistribute** tab.

Field	Description
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 and 65535.
MetricType	Specifies the metric type to associate with the route redistribution: <i>type1</i> or <i>type2</i> .
Subnets	Indicates whether subnetworks need to be advertised individually. Options available are: allow and suppress.
RoutePolicy	Specifies the name of preconfigured route policy to be used as the redistribution policy.

Applying an OSPF accept or redistribution policy

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. To apply a preconfigured OSPF accept policy, select the **OspfInFilterApply** check box.
5. To apply a preconfigured OSPF redistribution policy, select the **RedistributeApply** check box.
6. If you are applying OSPF redistribution policies, select the type of redistribution to apply from the available options in the **OspfApplyRedistribute** field.
7. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Applying Policy** tab.

Field	Definition
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. Available options are: <ul style="list-style-type: none">• none• direct• static• ip

Configuring the global IP forwarding next-hop status

Use this procedure to globally enable IP forwarding next-hop on the switch.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Globals** tab.
4. In the Forwarding Next Hop section, select the **AdminEnabled** box.
5. On the toolbar, click **Apply**.

Configuring an IP forwarding next-hop policy

Use this procedure to configure a policy for IP forwarding next-hop.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP tree, click the **Forwarding-nh Policy** tab.
4. Click **Insert**.
5. Complete the fields as required.
6. Click **Insert**.
7. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Forwarding-nh Policy** tab.

Field	Definition
Name	Specifies the name of the policy
MatchInetAddressType	Specifies the type of address used for matching.
MatchInetAddress	Specifies the source address to match.
MatchInetAddressMask	Specifies the length of the mask to match.
MatchPortMin	Specifies the minimum port number to match.
MatchPortMax	Specifies the maximum port number to match.
SetNextHopInetAddressType	Specifies the type of address used for the next-hop.
SetNextHopInetAddress	Specifies the next hop address to be used to forward the packet.
SetSecondNextHopInetAddressType	Specifies the type of address used for the secondary next-hop.
SetSecondNextHopInetAddress	Specifies the secondary next hop address to be used to forward the packet if the primary address (SetNextHopInetAddress) is unresolved but the secondary address is resolved.

Configuring an IP forwarding next-hop policy for an interface

Use this procedure to configure a policy for IP forwarding next-hop for an interface.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP tree, click the **Forwarding-nh Interface Policy** tab.
4. Click **Insert**.
5. Complete the fields as required.
6. Click **Insert**.
7. On the toolbar, click **Apply**.

Variable definitions

Variable	Value
Index	Specifies the VLAN ID.
PolicyName	Specifies the name of the policy associated with this interface.
Mode	Specifies the policy mode: drop or normal routing.
AdminStatus	Specifies enabled or disabled. Administratively enabling or disabling an entry will automatically apply the operation to all policy attachments for the specified interface. Only existing entries may be administratively disabled.

Chapter 24: DHCP relay configuration using ACLI

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

! Important:

DHCP relay uses a hardware resource that is shared by switch Quality of Service applications. When DHCP relay is enabled globally, the Quality of Service filter manager will not be able to use precedence 11 for configurations. For the filter manager to be able to use this resource, DHCP relay must be disabled for the entire unit or stack.

Prerequisites

- 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 to the destination DHCP server is available on the switch.

DHCP relay configuration procedures

To configure DHCP relay, perform the following steps:

1. Ensure that DHCP relay is enabled globally. (DHCP relay is enabled by default.)
2. Configure the DHCP relay forwarding path, specifying the VLAN IP as the DHCP relay agent and the remote DHCP server as the destination.
3. Enable DHCP for the specific VLAN.

Configuring global DHCP relay status

Use this procedure to configure the global DHCP relay status. DHCP relay is enabled by default.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip dhcp-relay
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables DHCP relay.

Displaying the global DHCP relay status

Use this procedure to display the current DHCP relay status for the switch.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip dhcp-relay
```

Specifying a local DHCP relay agent and remote DHCP server

Use this procedure to specify a 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

- Enable IP routing and configure an IP address on the VLAN to configure as a DHCP relay agent.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, 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 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}]	<p>Specifies the mode for DHCP relay.</p> <ul style="list-style-type: none"> • BootP only • BootP and DHCP • DHCP only <p>If you do not specify a mode, the default DHCP and BootP is used.</p>

Configuring DHCP relay maximum frame

Use this procedure to set the DHCP Relay maximum frame:

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[default] ip dhcp-relay [max-frame]
```

Displaying the DHCP relay configuration

Use this procedure to display the current DHCP relay agent configuration.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip dhcp-relay fwd-path [vrf {vrfName}] [vrfids {vrf_ids}]
```

Variable definitions

The following table describes the parameters for the `show ip dhcp-relay fwd-path` command.

Variable	Value
[vrf {vrfName}]	Specifies the VRF instance identified by the alphanumeric identifier (name) for which to display DHCP relay fwd-path information.
[vrfids {vrf_ids}]	Specifies the VRF instance(s) identified by the numerical ID(s) for which to display DHCP relay fwd-path information.

Job aid

The following table shows the field descriptions for the `show ip dhcp-relay fwd-path` command.

Field	Description
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 DHCP relay status and parameters on a VLAN

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.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, 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 command variables.

Variable	Value
[no]	Disables DHCP relay on the specified VLAN.
[broadcast]	Enables the broadcast of DHCP reply packets to the DHCP clients on this VLAN interface.
min-sec <min-sec>	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: <ul style="list-style-type: none"> • bootp - Supports BootP only • dhcp - Supports DHCP only • bootp_dhcp - Supports both BootP and DHCP
[option82]	Enables option 82 for DHCP Relay.

Configuring DHCP relay option 82 globally

Use this procedure to enable or disable DHCP Relay option 82 globally.

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip dhcp-relay option 82
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables DHCP relay globally.
[option 82]	Enables option 82 for DHCP Relay.

Configuring DHCP relay option 82 on a port

Use the following procedure to configure the subscriber ID for DHCP Relay option 82.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip dhcp-relay option82 <subscriber-id>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified subscriber ID.
<subscriber-id>	Specifies the subscriber ID for DHCP Relay option 82.

Displaying the DHCP relay configuration for a VLAN

Use this procedure to display the current DHCP relay parameters configured for a VLAN.

Procedure steps

1. Log on to the Privileged EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show vlan dhcp-relay [<vid>]
```

Variable definitions

The following table describes the 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
lflindex	Indicates the VLAN interface index.

Field	Description
MIN_SEC	Indicates the minimum time, in seconds, to wait between receiving a DHCP packet and forwarding the DHCP packet to the destination device. A value of zero indicates forwarding is done immediately without delay.
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 whether Option 82 is enabled or disabled on the VLAN.

Displaying DHCP relay counters

Use this procedure to display the current DHCP relay counters. This includes the number of requests and the number of replies.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, 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

Use this procedure to clear the DHCP relay counters for a VLAN.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip dhcp-relay clear-counters
```


Chapter 25: DHCP relay configuration using Enterprise Device Manager

The following chapter describe the procedures you can use to configure DHCP relay using Enterprise Device Manager (EDM).

Prerequisites

- 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 to the destination DHCP server is available on the switch.

Configuring DHCP relay using EDM

Use the following procedure to configure DHCP relay using EDM.

1. Enable DHCP relay globally.
2. Configure the DHCP relay forwarding path by specifying the VLAN IP as the DHCP relay agent and the remote DHCP server as the destination.
3. Enable DHCP relay on the VLAN.

Configuring global DHCP Relay status and parameters

Use the following procedure to configure the global DHCP Relay status and parameters.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay Globals** tab.
4. To enable DHCP Relay globally, select the **DhcpForwardingEnabled** box.
5. Configure the other parameters as required.
6. On the toolbar, click **Apply**.
7. On the toolbar, you can click **Refresh** to verify the global DHCP Relay configuration.

Variable definitions

The following table describes the fields of the **DHCP Relay Globals** tab.

Variable	Value
DhcpForwardingEnabled	Enables or disables the global DHCP forwarding status on the switch.
DhcpForwardingOption82 Enabled	Enables or disables option 82 on the switch.
DhcpForwardingMaxFrameLength	Configures the DHCP Forwarding frame length on the switch.

Configuring a DHCP Relay forwarding path using EDM

Use the following procedure to configure a DHCP Relay forwarding path.

Prerequisites

- Enable IP routing and DHCP Relay globally.
- Enable IP routing and configure an IP address on the VLAN to be set as the DHCP relay agent.
- Ensure that a route to the destination DHCP server is available on the switch.

Procedure steps

1. From the navigation DHCP Relay, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay** tab.
4. In the toolbar, click **Insert**.
5. Type the IP address of the local VLAN to serve as the DHCP relay agent in the **AgentAddr** field.
6. Type the remote DHCP Server IP address in the **ServerAddr** field.
7. Select the **Enable** check box if not selected.
8. Choose the desired DHCP relay mode in the **Mode** box.
9. Click **Insert**.
10. On the toolbar, you can click **Refresh** to verify the DHCP relay forwarding path configuration.

Variable definitions

The following table describes the fields of the **DHCP Relay** tab.

Field	Description
AgentAddr	Indicates the IP address of the local VLAN serving as the DHCP relay agent.
ServerAddr	Indicates 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 parameters on a VLAN using EDM

Use the following procedure to configure the DHCP relay parameters on a VLAN.

Procedure steps

1. From the navigation pane, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the **Basic** tab, select the VLAN for which DHCP relay is to be configured.
4. In the toolbar, click **IP**.
5. In the work area, click the **DHCP** tab.
6. Configure the parameters as required.
7. On the toolbar, click **Apply**.
8. On the toolbar, you can click **Refresh** to verify the VLAN DHCP configuration.

Variable definitions

The following table describes the fields of the **DHCP** tab.

Field	Description
Enable	Specifies whether DHCP relay is enabled or disabled.
MinSec	Specifies 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.
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.
Option82Enabled	Enables or disables option 82.
ClearCounters	Specifies to clear the DHCP relay counters for the VLAN.
CounterClearTime	Specifies the last time the counter values in this entry were reset to 0.

Configuring DHCP Relay option 82 on a VLAN

Use the following procedure to configure the option 82 parameter for DHCP Relay on a VLAN.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP Routing tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay-VLAN** tab.
4. To enable option 82, double-click under the **VlanDhcpOption82Enabled** field in the row containing the desired VLAN ID and select **true**.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** to verify the VLAN DHCP relay Option 82 configuration.

Variable definitions

The following table describes the fields of the **DHCP Relay-VLAN** tab.

Variable	Value
VlanDhcpOption82Enabled	Enables or disables option 82 on the specified VLAN.

Configuring DHCP Relay option 82 on a port

Use the following procedure to configure the option 82 parameter for DHCP Relay on a port.

Procedure steps

1. From the navigation pane, double-click **IP Routing**.
2. In the IP Routing tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay-port** tab.
4. In the **Switch/Stack/Ports** field, click the ellipsis (...) and select the ports to configure.
5. Under the **PortDhcpOption82SubscriberId** field, enter a subscriber ID.

6. Click **Apply Selection**.
7. On the toolbar, you can click **Refresh** to verify the port DHCP relay Option 82 configuration.

Variable definitions

The following table describes the fields of the **DHCP Relay-port** tab.

Field	Description
PortDhcpOption82SubscriberId	Specifies a subscriber ID for option 82 on the specified port.

Chapter 26: UDP broadcast forwarding configuration 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. This section describes how to configure UDP broadcast forwarding using ACLI.

The UDP broadcast forwarding feature cannot be enabled or disabled 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.

Prerequisites

- 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 to the destination address is available on the switch.

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 table entries that identify the protocols associated with specific UDP ports that you want to forward.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
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

1. Log on to the User Exec mode in ACLI.
2. At the command prompt, enter the following command:

```
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

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
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 of the UDP forwarding.
<list_name>	Specifies the name of the UDP forwarding list being created. (maximum 15 characters). The list name cannot be changed after first entering a port in the list.

Applying a UDP forwarding list to a VLAN

Use this procedure to associate a UDP forwarding list to 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

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip forward-protocol udp [vlan <vid>] [portfwddlist
<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 <code>interface vlan</code> 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 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 mask of the interface to which the list is attached is used. (See Note 1.)
<max_ttl>	Specifies the time 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.)
<p>Note 1: If you specify <code>maxttl</code> and/or <code>broadcastmask</code> values with no <code>portfwddlist</code> specified, then the switch saves the settings for this interface. When a <code>portfwddlist</code> is subsequently attached to this interface, and the <code>maxttl</code> and/or <code>broadcastmask</code> value are not defined, the saved parameters are automatically attached to the list. But if when specifying the <code>portfwddlist</code>, you also specify the <code>maxttl</code> and/or <code>broadcastmask</code>, your specified properties are used, regardless of any previous configurations.</p>	

Displaying the UDP broadcast forwarding configuration

Use this procedure to display the UDP broadcast forwarding configuration.

Procedure steps

1. Log on to the User Exec mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip forward-protocol udp [interface [vlan <1-4094>]]
[portfwdlist [<forward_list>]]
```

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.

Field	Description
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 portfwddlist` command.

Field	Description
LIST_ID	Specifies the specific 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

1. Log on to the Privileged Exec mode in ACLI.
2. At the command prompt, enter the following command:

```
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 27: 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.

This chapter describes the procedures that you can use to configure and manage UDP broadcast forwarding using Enterprise Device Manager (EDM).

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.
- Ensure that a route to the destination address is available on the switch.

Configuring UDP broadcast forwarding

Use the following procedure to configure UDP broadcast forwarding using EDM.

Procedure steps

1. Create UDP protocol entries that specify each UDP port and associated protocol that you want to forward.
2. Create UDP forwarding entries that specify the destination address for each UDP port that you want to forward.
3. Add UDP forwarding entries to a UDP forwarding list (you can create up to 128 UDP forwarding lists.)
4. Apply UDP forwarding lists to local VLAN interfaces.

Configuring UDP protocol table entries

Use the following procedure to create UDP table entries that identify the protocols associated with specific UDP ports that you want to forward.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.
- Ensure that a route to the destination address is available on the switch.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding 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** field.
6. Type the protocol name associated with the UDP port number in the **Name** field.
7. Click **Insert**.
8. On the toolbar, you can click **Refresh** to verify the UDP protocol table entries.

Variable definitions

The following table describes the fields of the **Protocols** tab.

Field	Description
PortNumber	Specifies the UDP port number.
Name	Specifies the protocol name associated with the UDP port.

Configuring UDP forwarding entries

Use the following procedure to configure individual UDP forwarding list entries, which associate UDP forwarding ports with destination IP addresses.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.
- Ensure that a route to the destination address is available on the switch.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding work area, click the **Forwarding** tab.
4. On the toolbar, click **Insert**.
5. Using the provided fields, specify a destination address for a selected port.
6. Click **Insert**.
7. On the toolbar, you can click **Refresh** to verify the UDP forwarding entry configuration.

Variable definitions

The following table describes the fields of the **Forwardings** tab.

Field	Description
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

Use the following procedure to add the UDP port/destination forwarding list entries (configured in the Forwardings tab) to UDP forwarding lists. Each UDP forwarding list can contain multiple port/destination entries.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.
- Ensure that a route to the destination address is available on the switch.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding work area, click the **Forwarding Lists** tab.
4. On the toolbar, click **Insert**.
5. Assign a unique ID to the UDP forwarding list in the **Id** field.
6. Type a unique name for the UDP forwarding list in the **Name** field.
7. Click the **FwdIdList (...)** to select the desired port/destination pairs from the list.
8. Click **Insert**.
9. On the toolbar, you can click **Refresh** to verify the UDP forwarding list configuration.

Variable definitions

The following table describes the fields of the **Forwarding Lists** tab.

Field	Description
Id	Indicates the unique identifier assigned to the forwarding list.
Name	Indicates the name assigned to the forwarding list.

Field	Description
FwdIdList	Indicates the forwarding entry IDs associated with the port/server IP pairs created using the Forwardings tab.

Applying a UDP forwarding list to a VLAN

Use the following procedure to assign a UDP forwarding list to a VLAN and to configure the related UDP forwarding parameters for the VLAN.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.
- Ensure that a route to the destination address is available on the switch.

Procedure steps

1. From the navigation pane, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding work area, click the **Broadcast Interfaces** tab.
4. On the toolbar, click **Insert**.
5. Click the **LocalIfAddr** ellipsis (...) to select a VLAN IP address from the list.
6. Click the **UdpPortFwdListId** ellipsis (...) to select the desired UDP forwarding list to apply to the VLAN from those listed.
7. Type the maximum TTL value in the **MaxTtl** field.
8. Type the broadcast mask in the **BroadCastMask** field.
9. Click **Insert**.
10. On the toolbar, you can click **Refresh** to verify the UDP forwarding list to VLAN configuration.

Variable definitions

The following table describes the fields of the **Broadcast Interface** tab.

Field	Description
LocalIfAddr	Indicates the IP address of the local VLAN interface.
UdpPortFwdListId	Indicates 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. This is an integer value between 1 and 16.
NumRxPkts	Indicates the total number of UDP broadcast packets received by this local interface.
NumFwdPkts	Indicates the total number of UDP broadcast packets forwarded.
NumDropPktsDestUnreach	Indicates the total number of UDP broadcast packets dropped because the destination was unreachable.
NumDropPktsUnknownPort	Indicates 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 mask of the interface to which the list is attached is used.

Chapter 28: Directed broadcasts configuration using ACLI

This chapter describes 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

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
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

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip directed-broadcast [interface [vlan <1-4094>] [vrf  
{vrfName}] [vrfids {vrf_ids}]
```

Variable definitions

The following table describes the parameters for the `show ip directed-broadcast` command.

Variable	Value
<1-4094>	Specifies the VLAN ID.
[vrf {vrfName}]	Specifies the VRF instance identified by the alphanumeric identifier (name) for which to display IP directed-broadcast information.
[vrfids {vrf_ids}]	Specifies the VRF instances identified by the numerical IDs for which to display IP directed-broadcast information.

Configuring IP directed broadcasts for each VLAN

Use this procedure to configure IP directed broadcasts for each VLAN on the switch. By default, IP directed broadcasts are disabled.

Enabling IP directed broadcasts for each VLAN

Use this procedure to enable the IP directed broadcasts for each VLAN.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip directed-broadcast [enable]
```

Disabling IP directed broadcasts for each VLAN

Use this procedure to disable IP directed broadcasts for each VLAN.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
no ip directed-broadcast [enable]
```

Setting IP directed broadcasts for each VLAN to default

Use this procedure to set IP directed broadcasts to default.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
default ip directed-broadcast [enable]
```


Chapter 29: Static ARP and Proxy ARP configuration using ACLI

This chapter describes the procedures you can use to configure Static Address Resolution Protocol (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 create and enable a static ARP entry.

Prerequisites

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

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip arp <A.B.C.D> <aa:bb:cc:dd:ee:ff> <unit / port> [vid  
<1-4094>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified ARP entry.
<A.B.C.D>	Specifies the IP address of the device to configure as a static ARP entry.
<aa:bb:cc:dd:ee:ff>	Specifies the MAC address of the device to configure as a static ARP entry.
<unit / port>	Specifies the unit and port number to which you add the static ARP entry.
vid <1 - 4094>	Specifies the VLAN ID to which you add the static ARP entry.

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:

```
5530-24TFD(config)# ip arp 10.1.1.23 00:00:11:43:54:23 1/48 vid 1
```

Configuration example: Deleting a static ARP entry

The following is an example of deleting a static ARP entry:

```
5530-24TFD(config)# no ip arp 172.2.2.13
```

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.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show arp-table
```

OR

```
show arp [static | dynamic] [<ip-addr> | {-s <subnet>
<mask>}] [summary] [vrf {vrfName}] [vrfids {vrf_ids}]
```

OR

```
show ip arp [static | dynamic] [<ip-addr> | {-s <subnet>
<mask>}] [summary] [vrf {vrfName}] [vrfids {vrf_ids}]
```

The **show ip arp** command is invalid if the switch is not in Layer 3 mode.

Variable definitions

The following table describes the command variables.

Variable	Value
<ip-addr>	Specifies the IP address of the ARP entry to display.
-s <subnet> <mask>	Displays ARP entries for the specified subnet only.
static	Displays all configured static entries, including those without a valid route.
vrf {vrfName}	Specifies the VRF instance identified by the alphanumeric identifier (name) for which to display route information.
vrfids {vrf_ids}	Specifies the VRF instances identified by the numerical IDs for which to display route information.

Job aid

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

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.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. To configure a global timeout for ARP entries, enter the following command:

```
ip arp timeout <timeout>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<timeout>	Specifies the amount of time in minutes before an ARP entry ages out. The range is 5-360. The default value is 360 minutes.

Configuration example: Changing the global ARP timeout

The following is an example of setting a new default aging time:

```
5530-24TFD(config)# ip arp timeout 180
```

The new setting can be confirmed by using the `show ip routing` command.

Clearing the ARP cache

Use this procedure to clear the cache of ARP entries.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. To clear the ARP cache, enter the following command:

```
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.

Prerequisite

- 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

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. To configure proxy ARP status, enter the following command:

```
[default] [no] ip arp-proxy enable
```

Variable definitions

The following table describes the 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

1. Log on to the User EXEC mode in ACLI.
2. To display proxy ARP status for a VLAN, enter the following command:

```
show ip arp-proxy interface [vlan <vid>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vid>	Specifies the ID of the VLAN to display. The range is 1-4094.

Job aid

The following table shows the field descriptions for the `show ip arp-proxy interfaces` command.

Field	Description
Vlan	Identifies a VLAN.
Proxy ARP status	Specifies the status of Proxy ARP on the VLAN.

Chapter 30: Static ARP and Proxy ARP configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure static ARP and proxy ARP using Enterprise Device Manager (EDM).

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a Proxy ARP interface.

Configuring static ARP entries

Use the following procedure to configure static ARP entries for the switch.

Prerequisites

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

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IP**.
3. In the work area, click the **ARP** tab.
4. In the toolbar, click **Insert**.

5. Click the **Port in VLAN** button in the **Interface** field to select the VLAN to which the static ARP entry is being added.
6. Select the port for this ARP entry in the **VLAN** dialog box.
7. Type the IP address for the ARP entry in the **IPAddress** field.
8. Type the MAC address for the ARP entry in the **MacAddress** field.
9. Click **Insert**.

Variable definitions

The following table describes the fields of the **Insert ARP** tab.

Field	Description
Interface	Specifies the VLAN name and port to which you add the static ARP entry.
MacAddress	Specifies the MAC address of the device you set as a static ARP entry.
IpAddress	Specifies the IP address of the device you set as a static ARP entry.
Type	Specifies the type of ARP entry: static, dynamic, local, or broadcast.

Configuring Proxy ARP

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.

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

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IP**.
3. In the work area, click the **ARP Interfaces** tab.

! **Important:**

EDM does not display the ARP Interfaces tab if you have not enabled routing on the switch.

4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select **enable** or **disable** from the drop-down list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **ARP Interfaces** tab.

Field	Description
IfIndex	Specifies a configured switch interface.
DoProxy	Enables or disables proxy ARP on the interface.
DoResp	Enables or disables the sending of ARP responses on the specified interface.

Chapter 31: 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

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip blocking-mode {full | none}
```

Variable definitions

The following table describes the command variables.

Variable	Value
full	Select this parameter to configure IP blocking to full. This never allows a duplicate IP address in a stack.
none	Select this parameter to configure IP blocking to none. This 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. Log on to User EXEC mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip blocking
```


Chapter 32: VRRP configuration using ACLI

The Virtual Router Redundancy Protocol (VRRP) eliminates 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.

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

Use this procedure to configure the global VRRP status on the switch.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] router vrrp enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Globally disables VRRP on the switch.

Assigning an IP address to a virtual router ID

Use this procedure to associate an IP address with a virtual router ID.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip vrrp address <vr_id> <ip_address>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the IP address from the virtual router ID.
<vr_id>	Specifies the virtual router to configure. This is a value between 1 and 255.
<ip_address>	Represents the address to associate with the virtual router ID.

Configuring the router priority for a virtual router ID

Use this procedure to assign a priority to the router for a specified virtual router ID.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip vrrp <vr_id> priority <priority_value>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router ID to configure for the router priority.
<priority_value>	Specifies the priority to assign to the router for the specified virtual router ID. This is a value between 1 and 255.

Configuring the status of the virtual router

Use this procedure to enable the virtual router.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip vrrp <vr_id> enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
[no]	Disables the virtual router.

Configuring a backup master

Use this procedure to configure the VRRP backup master functionality. Enable backup master on both the master and backup routers.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip vrrp <vr_id> backup-master enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables the VRRP backup master functionality.
<vr_id>	Specifies the virtual router to configure.

Configuring the critical IP address

Use this procedure to set the critical IP address on the router.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip vrrp <vr_id> critical-ip-addr <ip_address>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
<ip_address>	Specifies the critical IP address to use.
[no]	Removes the configured critical IP.

Configuring the VRRP critical IP status

Use this procedure to configure the status of the VRRP critical IP functionality.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip vrrp <vr_id> critical-ip enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
[no]	Disables the VRRP critical IP functionality.

Configuring the VRRP holddown timer

Use this procedure to set the VRRP holddown timer.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip vrrp <vr_id> holddown-timer <timer_value>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
<timer_value>	Specifies the holddown timer value. This is a value in seconds between 0 and 21600.

Configuring VRRP holddown action

Use this procedure to define the action this interface takes when a holddown timer threshold is reached.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI
2. At the command prompt, enter the following command:

```
ip vrrp <vr_id> action {none | preempt}
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
{none preempt}	If you select <i>none</i> as the action, the device takes no action when a holddown timer threshold is reached. If you select <i>preempt</i> as the action, the holddown timer is cancelled.

Configuring the VRRP advertisement interval

Use this procedure to configure the VRRP advertisement interval.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip vrrp <vr_id> adver-int <interval>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
<interval>	Specifies the advertisement interval in seconds. This is an integer value between 1 and 255.

Configuring the fast advertisement interval

Use this procedure to set the interval used in the VRRP fast advertisement functionality.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip vrrp <vr_id> fast-adv-int <interval>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
<interval>	Specifies the fast advertisement interval. This is a value in milliseconds between 200 and 1000.

Configuring fast advertisement status

Use this procedure to enable the VRRP fast advertisement functionality.

Procedure steps

1. Log on to the Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip vrrp <vr_id> fast-adv enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router to configure.
[no]	Disables the VRRP fast advertisement functionality.

Configuring ICMP echo replies

Use this procedure to configure ICMP echo replies from VRRP associated addresses.

Procedure steps

1. Log on to the VRRP Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ping-virtual-address enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables ICMP echo replies from VRRP associated addresses.

Enabling VRRP traps

Use this procedure to enable the sending of SNMP notifications after virtual router state changes.

Procedure steps

1. Log on to the VRRP Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
send-trap enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables the sending of SNMP notifications after virtual router state changes.

Displaying VRRP configuration and statistics

Use this procedure to display VRRP configuration information and statistics. If you do not specify any parameters, the device displays basic global configuration information.

Procedure steps

1. Log on to the User EXEC mode in ACLI.
2. At the command prompt, enter the following command to display global VRRP properties:

```
show ip vrrp
```

3. At the command prompt, enter the following command to display the VRRP virtual IP address configuration for a VLAN:

```
show ip vrrp address [addr <A.B.C.D>] [vrid <1-255>]
[interface [addr <A.B.C.D>] [vlan <1-4094>] [vrid <1-255>]]
```

4. At the command prompt, enter the following command to display detailed VRRP interface configuration information:

```
show ip vrrp interface [verbose] [vlan <1-4094>] [vrid
<1-255>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
addr <A.B.C.D>	Displays VRRP configurations associated with the specified IP address.
interface [addr <A.B.C.D>] [vlan <1-4094>] [vrid <1-255>]	Displays VRRP configurations associated with the specified interface.

Variable	Value
vrid <1-255>	Displays VRRP configurations associated with the specified virtual router ID.
vlan <1-4094>	Displays VRRP configurations associated with the specified VLAN.
verbose	Displays additional VRRP configuration information.

Chapter 33: VRRP configuration examples using ACLI

This section provides configuration examples to configure VRRP on the Avaya Ethernet Routing Switch 5000 Series .

Configuring normal VRRP operation

The following configuration example shows how to provide VRRP service for two edge host locations.

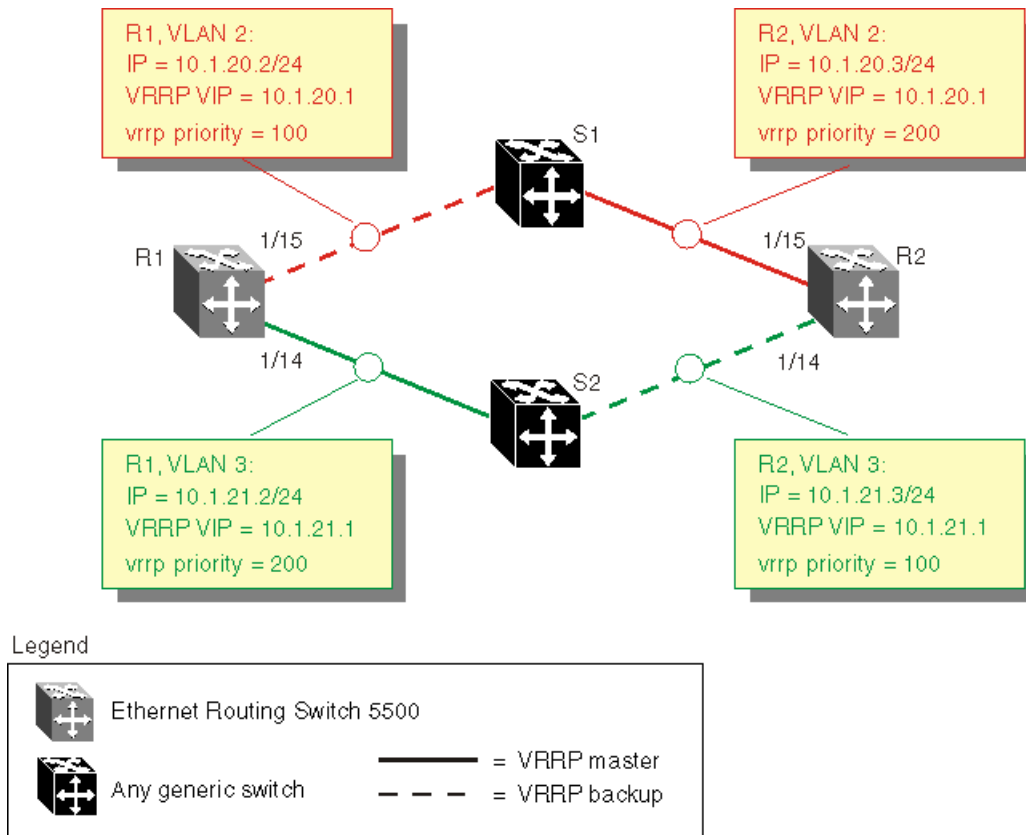


Figure 46: VRRP example topology

In this example, the switches have the following duties:

- R1 is the VRRP master for S2
- R2 is the VRRP master for S1

In this example, the administrator enables VRRP with OSPF as the routing protocol on R1 and R2.

The administrator uses VRRP priority setting 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 becomes the tie breaker. Therefore, it is very important to set the correct VRRP priority. The administrator also enables VRRP fast advertisement in this example to allow for fast failover detection.

The following procedure describes the steps necessary to reproduce the example above:

1. Configure VLAN 2 on router R1.
 - a. Create VLAN 2 on router R1.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

- b. Configure the ports for VLAN 2 on R1.

```
5530-24TFD# config terminal
5530-24TFD(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.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.2
255.255.255.0
```

- d. Configure an OSPF interface for VLAN 2.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.2
```

- e. Configure VRRP on VLAN 2.

The administrator adds VRRP VIP address of 10.1.20.1 to VLAN 2 using a VRID of 1.

*** Note:**

The administrator does not configure VRRP priority; it is left at the factory default of 100. Instead, the priority setting on router R2 will be 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. If you want fast VRRP advertisement, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip vrrp address 1 10.1.20.1
5530-24TFD(config-if)# ip vrrp 1 enable
```

2. Configure VLAN 3 on router R1.

- a. Configure VLAN 3 on router R1 using spanning tree group 1.

```
5530-24TFD# config terminal
5530-24TFD# vlan create 3 type port
```

- b. Configure the ports for VLAN 3 on R1.

```
5530-24TFD# config terminal
5530-24TFD(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.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 3
```

```
5530-24TFD(config)# ip address 10.1.21.2
255.255.255.0
```

d. Configure an OSPF interface for VLAN 3.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(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.

*** 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. If you want fast VRRP advertisement, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip vrrp address 2 10.1.21.1
5530-24TFD(config-if)# ip vrrp 2 priority 200
5530-24TFD(config-if)# ip vrrp 2 enable
```

3. Configure VLAN 2 on router R2.

a. Create VLAN 2 on router R2.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b. Configure the ports for VLAN 2 on R2.

```
5530-24TFD# config terminal
5530-24TFD(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.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.3
255.255.255.0
```

d. Configure an OSPF interface for VLAN 2.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(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.

*** 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. If you want fast VRRP advertisement, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip vrrp address 1 10.1.20.1
5530-24TFD(config-if)# ip vrrp 1 enable
5530-24TFD(config-if)# ip vrrp 1 priority 200
```

4. Configure VLAN 3 on router R2.

a. Configure VLAN 3 on router R2.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 3 type port
```

b. Configure the ports for VLAN 3 on R1.

```
5530-24TFD# config terminal
5530-24TFD(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.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip address 10.1.21.3 255.255.255.0
```

d. Configure an OSPF interface for VLAN 3.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(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.

*** 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. If you want fast VRRP advertisement, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 3
```

```
5530-24TFD(config-if)# ip vrrp address 2 10.1.21.1
5530-24TFD(config-if)# ip vrrp 2 enable
```

After you complete the VRRP configuration, use the **show ip vrrp** and **show ip vrrp interface verbose** commands to display VRRP configuration information and statistics.

Configuration command listing

This following list is a complete sequence of the commands used in this configuration:

1. VLAN Configuration for Router R1:

```
configure terminal
vlan create 2 type port
vlan members remove 2 1/1-1/14,2/1-2/8,3/1-3/8
vlan members add 2 1/15
interface vlan 2
ip address 10.1.20.2 255.255.255.0
router ospf enable
router ospf
network 10.1.20.2
router vrrp ena
interface vlan 2
ip vrrp address 1 10.1.20.1
ip vrrp 1 enable
vlan create 3 type port
interface vlan 3
ip address 10.1.21.2 255.255.255.0
router ospf enable
router ospf
network 10.1.21.2
router vrrp ena
interface vlan 3
ip vrrp address 2 10.1.21.1
ip vrrp 2 priority 200
ip vrrp 2 enable
```

2. VLAN Configuration for Router R2:

```
configure terminal
vlan create 2 type port
vlan members remove 2 1/1-1/14,2/1-2/8,3/1-3/8
vlan members add 2 1/15
interface vlan 2
ip address 10.1.20.3 255.255.255.0
router ospf enable
router ospf
network 10.1.20.3
router vrrp ena
interface vlan 2
ip vrrp address 1 10.1.20.1
ip vrrp 1 priority 200
ip vrrp 1 enable
vlan create 3 type port
vlan members remove 3 1/1-1/14,1/15,2/1-2/8,3/1-3/8
vlan members add 3 1/14 interface vlan 3
ip address 10.1.21.3 255.255.255.0
router ospf enable
router ospf
network 10.1.21.3
router vrrp ena
interface vlan 3
```

```
ip vrrp address 2 10.1.21.1
ip vrrp 2 enable
```

Configuring VRRP with SMLT

This configuration example shows how you can provide high availability for a Layer 2 edge switch feeding into a Layer 3 core. As demonstrated below, both R1 and R2 switches are configured with a port-based VLAN (VLAN 2) with SMLT and VRRP set to enabled. This topology provides failover protection and load-balancing.

The Avaya Ethernet Routing Switch 5000 Series (E1), running in Layer 2 mode, is configured with one port-based VLAN and one MultiLink Trunking (MLT) group for the aggregate uplink ports. The Avaya Ethernet Routing Switch 5000 Series switches (R1 and R2) are configured with backup master enabled so that both switches can reply to ARP.

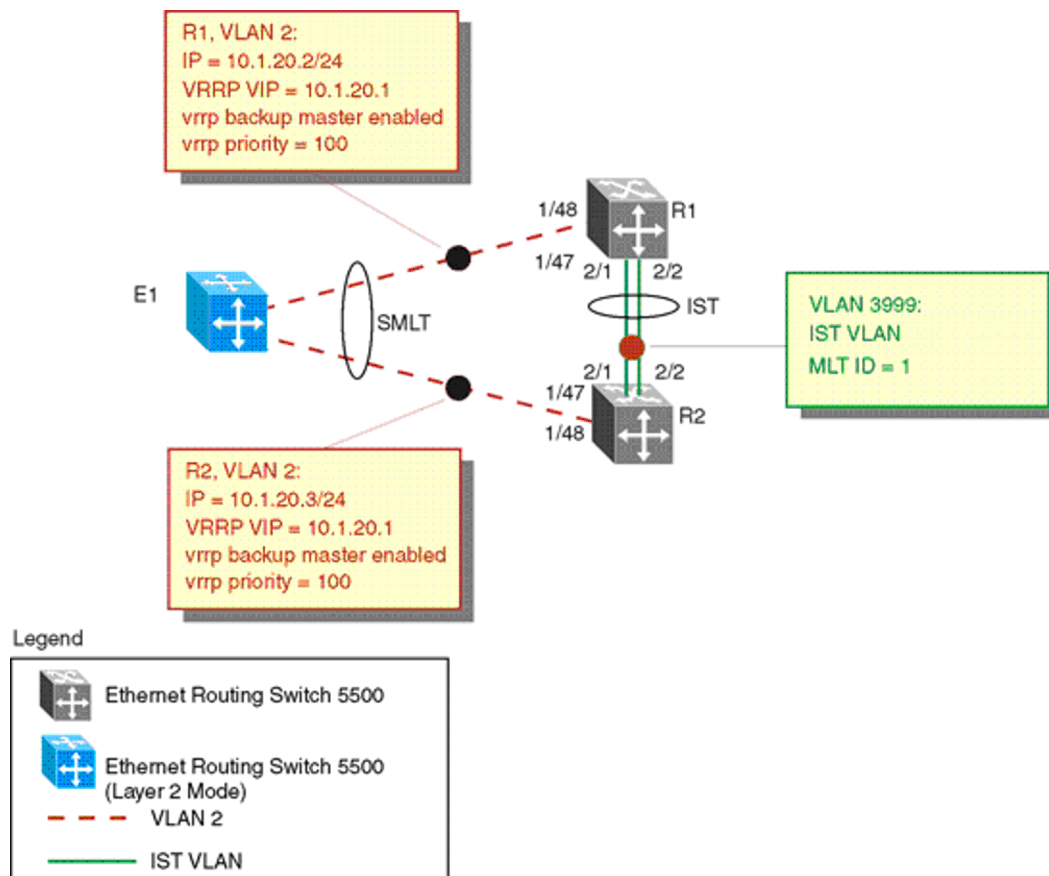


Figure 47: VRRP with SMLT configuration

Use the following procedure to recreate the illustrated topology:

1. Configure the IST VLAN on R1.

a. Configure IST VLAN 3999 on R1:

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 3999 type port
5530-24TFD(config)# interface vlan 3999
5530-24TFD(config-if)# ip address 2.1.1.1
255.255.255.0
```

b. Configure IST MLT on R1:

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 1 member 2/1-2/2
5530-24TFD(config)# vlan port 2/1-2/2 tagging enable
5530-24TFD(config)# mlt 1 enable
```

c. Configure the IST and add the IST to VLAN 3999:

```
5530-24TFD# config terminal
5530-24TFD(config)# interface mlt 1
5530-24TFD(config-if)# ist enable peer-ip 2.1.1.2
vlan 3999
```

2. Configure VRRP and SMLT for access VLAN to E1.

a. Configure VLAN 2 on R1:

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b. Create IP address for VLAN 2:

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.2
255.255.255.0
```

c. Configure the access port for VLAN 2 on R1 and add VLAN 2 to the IST and SMLT groups:

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 2 1/48,1/47,2/
1,2/2
```

*** Note:**

2/1 and 2/2 are IST ports. 1/48,1/47 are SMLT ports.

d. Create SMLT on R1:

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 2 member 1/47,1/48
5530-24TFD(config)# mlt 2 enable
5530-24TFD(config)# interface mlt 2
5530-24TFD(config-if)# smlt 1
```

e. Enable OSPF interface on VLAN 2 of R1:

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
```

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.2
```

f. Configure VRRP VIP address for VLAN2 of R1:

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config)# ip vrrp address 1 10.1.20.1
5530-24TFD(config)# ip vrrp 1 enable
5530-24TFD(config)# ip vrrp 1 backup-master 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. If you want fast VRRP advertisement, enable fast advertisement.

3. Configure the IST VLAN for router R2.

a. Configure IST VLAN 3999 on R2:

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 3999 type port
5530-24TFD(config)# interface vlan 3999
5530-24TFD(config-if)# ip address 2.1.1.2
255.255.255.0
```

b. Configure IST MLT on R2:

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 1 member 2/1-2/2
5530-24TFD(config)# mlt 1 enable
5530-24TFD(config)# vlan port 2/1-2/2 tagging enable
```

c. Configure an IST peer for R2 and add the IST to VLAN 3999:

```
5530-24TFD# config terminal
5530-24TFD(config)# interface mlt 1
5530-24TFD(config-if)# ist enable peer-ip 2.1.1.2
vlan 3999
```

4. Configure VRRP and SMLT for VLAN access to E1:

a. Configure VLAN 2 on R2:

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b. Create an IP address for VLAN 2:

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.3
255.255.255.0
```

- c. Configure the access port for VLAN 2 on R2 and add VLAN 2 to the IST and SMLT groups:

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 2 1/47, 1/48,
2/1. 2/2
```

*** Note:**

1/47 and 1/48 are SMLT ports. 2/1 and 2/2 are IST ports.

- d. Create SMLT on R2:

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 2 member 1/47, 1/48
5530-24TFD(config)# mlt 2 enable
5530-24TFD(config)# interface mlt 2
5530-24TFD(config-if)# smlt 1
```

- e. Enable OSPF interface for VLAN 2 on R2:

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.3
```

- f. Configure VRRP VIP address for VLAN 2 on R2:

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip vrrp address 1 10.1.20.1
5530-24TFD(config-if)# ip vrrp 1 enable
5530-24TFD(config-if)# ip vrrp 1 backup-master 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. If you want fast VRRP advertisement, enable fast advertisement.

Configuration command listing

This following list is a complete sequence of the commands used in this configuration:

1. Configuration for R1:

```
#MLT CONFIGURATION #
configure terminal
mlt 1 member 2/1-2/2
mlt 1 ena
vlan port 2/2-2/2 tagging enable
interface mlt 1
ist enable peer-ip 2.1.1.2
vlan 3999
```

```

mlt 2 member 1/48,1/47
mlt 2 enable
interface mlt 2
smlt 1

#VLAN CONFIGURATION #
configure terminal
vlan members remove 1 1/47-1/48,2/1-2/2
vlan create 2 type port
vlan members remove 2 1/1-1/46,2/3-2/8,3/1-3/8
vlan members add 2 1/47-1/48,2/1-2/2
interface vlan 2
ip address 10.1.20.2 255.255.255.0
router ospf enable
router ospf
network 10.1.20.2
router vrrp ena
interface vlan 2
ip vrrp address 1 10.1.20.1
ip vrrp 1 enable
ip vrrp 1 backup-master enable
vlan create 3999 type portvlan
members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8
vlan members add 3999 2/1-2/2
interface vlan 3999
ip address 2.1.1.1
255.255.255.0
#PORT CONFIGURATION #

#PHASE II #
configure terminal
mlt spanning-tree 1 stp all learning disable
mlt spanning-tree 2 stp all learning disable

```

2. Configuration for R2:

```

#MLT CONFIGURATION #

configure terminal
mlt 1 member 2/1-2/2
mlt 1 enable
vlan port 2/2-2/2 tagging enable
interface mlt 1
ist enable peer-ip 2.1.1.1
vlan 3999
mlt 2 member 1/48,1/47
mlt 2 enable
interface mlt 2
smlt 1

#VLAN CONFIGURATION #

configure terminal
vlan members remove 1 1/47-1/48,2/1-2/2
vlan create 2 type port
vlan members remove 2 1/1-1/46,2/3-2/8,3/1-3/8
vlan members add 2 1/47-1/48,2/1-2/2
interface vlan 2
ip address 10.1.20.3 255.255.255.0
router ospf enable
router ospf network 10.1.20.2
router vrrp ena
interface vlan 2
ip vrrp address 1 10.1.20.1
ip vrrp 1 enable
ip vrrp 1 backup-master enable

```

```
vlan create 3999 type portvlan
members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8
vlan members add 3999 2/1-2/2
interface vlan 3999
ip address 2.1.1.2 255.255.255.0

#PORT CONFIGURATION #

# PHASE II #
configure terminal
mlt spanning-tree 1 stp all learning disable
mlt spanning-tree 2 stp all learning disable
```

Configuring VRRP with SLT

The following illustration and configuration file examples demonstrate a VRRP configuration with SLT.

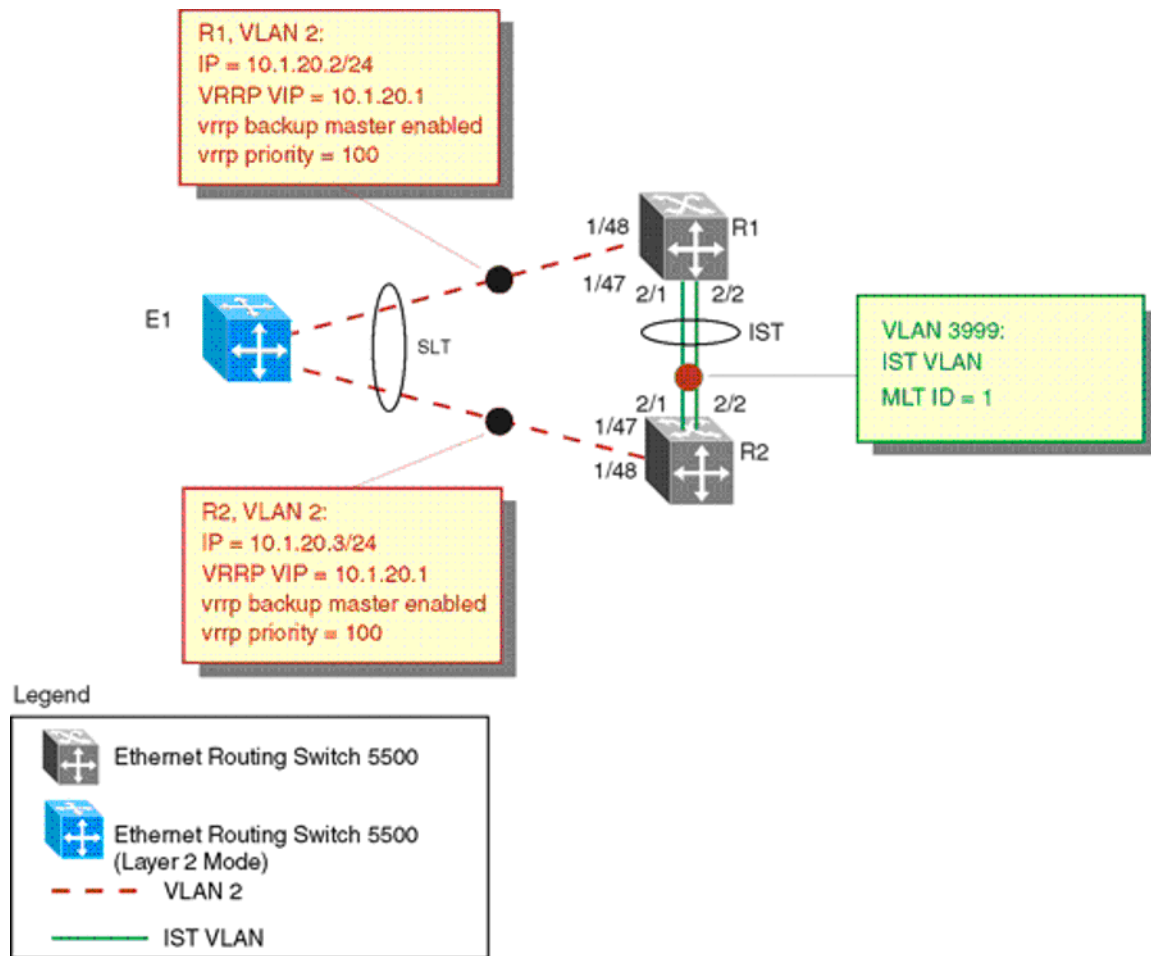


Figure 48: VRRP with SLT configuration

Use the following commands to recreate the above configuration:

1. Configuration for R1:

```
#MLT CONFIGURATION #
configure terminal
mlt 1 member 2/1-2/2
mlt 1 enable
vlan port 2/2-2/2 tagging enable
interface mlt 1
ist enable peer-ip 2.1.1.2
vlan 3999
interface fast-Ethernet 1/48
smlt 1

#VLAN CONFIGURATION #
configure terminal
vlan members remove 1 1/48,2/1-2/2
vlan create 2 type port
vlan members remove 2 1/1-1/47,2/3-2/8,3/1-3/8
vlan members add 2 1/47,2/1-2/2
interface vlan 2
```

VRRP configuration examples using ACLI

```
ip address 10.1.20.2 255.255.255.0
router ospf enable
router ospf
network 10.1.20.2
router vrrp ena
interface vlan 2
ip vrrp address 1 10.1.20.1
ip vrrp 1 enable
ip vrrp 1 backup-master enable
vlan create 3999 type port
vlan members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8
vlan members add 3999 2/1-2/2
interface vlan 3999
ip address 2.1.1.1 255.255.255.0

#PORT CONFIGURATION #
# PHASE II #
configure terminal
mlt spanning-tree 1 stp all learning disable
interface fast-Ethernet 1/48 spanning-tree stp 1 learning disable
```

2. Configuration for R2:

```
#MLT CONFIGURATION #

configure terminal
mlt 1 member 2/1-2/2
mlt 1 enable
vlan port 2/2-2/2 tagging enable
interface mlt 1
ist enable peer-ip 2.1.1.1
vlan 3999
interface fast-Ethernet 1/48
smlt 1

#VLAN CONFIGURATION #
configure terminal
vlan members remove 1 1/48,2/1-2/2
vlan create 2 type port
vlan members remove 2 1/1-1/47,2/3-2/8,3/1-3/8
vlan members add 2 1/48,2/1-2/2
interface vlan 2
ip address 10.1.20.3 255.255.255.0
router ospf enable
router ospf
network 10.1.20.2
router vrrp ena
interface vlan 2
ip vrrp address 1 10.1.20.1
ip vrrp 1 enable
ip vrrp 1 backup-master enable
vlan create 3999 type port
vlan members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8
vlan members add 3999 2/1-2/2
interface vlan 3999
ip address 2.1.1.2 255.255.255.0

#PORT CONFIGURATION #
#PHASE II #
configure terminal
mlt spanning-tree 1 stp all learning disable interface
fast-Ethernet 1/48 spanning-tree stp 1 learning disable
```

Chapter 34: VRRP configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure VRRP using Enterprise Device Manager (EDM).

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.

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.

Configuring VRRP

Use the following procedure to enable VRRP on a VLAN.

Procedure 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 and enable the virtual router.

Configuring global VRRP status and properties

Use the following procedure to configure global VRRP settings.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **VRRP**.
3. In the work area, click the **Globals** tab.
4. Configure the global VRRP parameters as required.
5. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **VRRP—Globals** tab.

Field	Definition
Enabled	Indicates whether VRRP is globally enabled on the switch.
Version	Indicates the version of VRRP that the switch supports.
NotificationCntl	Indicates whether the VRRP-enabled router generates Simple Network Management Protocol (SNMP) traps based on VRRP events: <ul style="list-style-type: none">• Enabled - SNMP traps are sent.• Disabled - SNMP traps are not sent.
PingVirtualAddrEnabled	Indicates whether this device responds to pings directed to a virtual router IP address.

Assigning an IP address to a virtual router ID

Use the following procedure to associate an IP address with a virtual router ID on a switch interface.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **VRRP**.
3. In the work area, click the **Interface Address** tab.
4. In the toolbar, click **Insert**.
5. Using the provided fields, create the new interface.
6. Click **Insert**.

Variable definitions

The following table describes the fields of the **Interface Address** tab.

Field	Definition
IfIndex	Specifies the IP address of the interface on which to configure VRRP.
VrId	Specifies the virtual router ID to associate with this interface.
IpAddr	Specifies the IP address to associate with the virtual router ID.
Status	Specifies the status of the interface; active or inactive.

Configuring VRRP interface properties

Use the following procedure to configure VRRP interface properties.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **VRRP**.
3. In the work area, click the **Interfaces** tab.
4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select a parameter or value from the drop-down list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Interfaces** tab.

Field	Definition
Index	Specifies the index of the VRRP interface.
VrId	Specifies the virtual router ID.
PrimaryIpAddr	Specifies an IP address selected from the set of real interface addresses. VRRP advertisements are always sent using the primary IP address as the source of the IP packet.

Field	Definition
VirtualMacAddr	Specifies the virtual MAC address of the virtual router.
State	Specifies the current state of the virtual router. A virtual router can be in one of the following states: <ul style="list-style-type: none"> • initialize indicates the virtual router is waiting for a startup event. • backup indicates the virtual router is monitoring the availability of the master router. • master indicates the router is forwarding packets for IP addresses associated with the virtual router ID.
AdminState	Indicates the administrative status of the virtual router.
Priority	Indicates the priority to use for the virtual router master election process. This field is an integer value between 1 and 255. The priority value for the VRRP router that owns the IP addresses associated with the virtual router must be 255. The default priority value for VRRP routers that back up a virtual router is 100.
MasterIpAddr	Indicates the real (primary) IP address of the master router. This address is the IP address listed as the source in the VRRP advertisement last received by this virtual router.
AdvertisementInterval	Indicates the time interval, in seconds, between transmission of advertisement messages. Only the master router sends VRRP advertisements. This field is an integer value between 1 and 255. The default value is 1.
VirtualRouterUpTime	Indicates the amount of time this virtual router has spent out of the initialize state.
HolddownTimer	Specifies the amount of time (in seconds) to wait before preempting the current VRRP master. This field is an integer value between 0 and 21600.
HoldDownState	Specifies the holddown state of this VRRP interface.
HoldDownTimeRemaining	Specifies the amount of time (in seconds) left before the holddown timer expires.
Action	Triggers an action on this VRRP interface. Options available are: none (no action) or preemptHoldDownTimer.
CriticalIpAddrEnabled	Indicates whether the user-defined critical IP address is enabled. If the user-defined critical IP address is not enabled, a default critical IP address of 0.0.0.0 is used.
CriticalIpAddr	Specifies the IP address of the interface that will cause a shutdown event.
BackupMasterEnabled	Indicates whether the backup/master functionality is enabled on this interface.

Field	Definition
BackupMasterState	Indicates the state of the backup/master functionality.
FastAdvertisementEnabled	Indicates if the Faster Advertisement Interval should be used. The default value is false.
FastAdvertisementInterval	Specifies the fast advertisement interval, in milliseconds, between sending advertisement messages. This field is an integer value between 200 and 1000. The default value is 200.

Graphing VRRP interface information

Use the following procedure to display and graph VRRP interface statistical information.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **VRRP**.
3. In the work area, click the **Interfaces** tab.
4. In the table, select the desired interface.
5. In the toolbar, click **Graph**.
6. Using the provided fields, view and graph the VRRP statistical information.

Variable definitions

The following table describes the fields of the **VRRP Stats** tab.

Field	Definition
BecomeMaster	Specifies the total number of times that the state of this virtual router has transitioned to master.
AdvertiseRcvd	Specifies the total number of VRRP advertisements received by this virtual router.
AdvertiseIntervalErrors	Specifies the total number of VRRP advertisement packets received for which the advertisement interval is different than the one configured for the local virtual router.
IpTtlErrors	Specifies the total number of VRRP packets received by the virtual router with IP Time-To-Live (TTL) not equal to 255.
PriorityZeroPktsRcvd	Specifies the total number of VRRP packets received by the virtual router with a priority of 0.
PriorityZeroPktsSent	Specifies the total number of VRRP packets sent by the virtual router with a priority of 0.
InvalidTypePktsRcvd	Specifies the number of VRRP packets received by the virtual router with an invalid value in the type field.
AddressListErrors	Specifies the total number of packets received for which the address list does not match the locally configured list for the virtual router.
AuthFailures	Specifies the total number of VRRP packets received that do not pass the authentication check.
InvalidAuthType	Specifies the total number of packets received with an unknown authentication type.
AuthTypeMismatch	Specifies the total number of packets received with Auth Type not equal to the locally configured authentication method.
PacketLengthErrors	Specifies the total number of packets received with a packet length less than the length of the VRRP header.

Viewing general VRRP statistics

Use the following procedure to view and graph general VRRP statistics.

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.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **VRRP**.
3. In the work area, click the **Stats** tab.

Variable definitions

The following table describes the fields of the **Stats** tab.

Field	Definition
RouterChecksumErrors	Specifies the total number of VRRP packets received with an invalid VRRP checksum value.
RouterVersionErrors	Specifies the total number of VRRP packets received with an unknown or unsupported version number.
RouterVrIdErrors	Specifies the total number of VRRP packets received with an invalid VRID for this virtual router.

Chapter 35: IGMP snooping configuration using ACLI

This chapter describes the procedures you can use to configure IGMP snooping on a VLAN using ACLI.

IGMP snooping configuration procedures

To configure IGMP snooping, you must enable snooping on the VLAN.

All related configurations, listed below, are optional. You can use these procedures to suit the requirements of your network.

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>	
Global Configuration mode	

Command	Parameter
<code>ip pgmp</code>	<code>flush vlan <vid> <grp-member mrouter sender></code>
<code>vlan igmp <vid></code>	<code>[snooping {enable disable}]</code>
	<code>[proxy {enable}disable]</code>
	<code>[query-interval <query-int>] [robust-value <robust-val>]</code>
	<code>unknown-mcast-no-flood {enable disable}</code>
	<code>unknown-mcast-allow-flood <H.H.H> <mcast_ip_address></code>
	<code>{vl-members v2-members} [add remove] <portlist></code>
Privileged EXEC mode	
<code>show ip igmp</code>	<code>cache</code>
	<code>group [count] [group <A.B.C.D>] [member-subnet <A.B.C.D>/<0-32>]</code>
	<code>interface [vlan <vid>]</code>
	<code>snooping</code>
<code>show vlan igmp</code>	<code>unknown-mcast-allow-flood</code>
	<code>unknown-mcast-no-flood</code>
	<code><vid></code>
<code>show vlan multicast membership</code>	<code><vid></code>

From Release 6.2 onwards there is a method you can use to associate IGMP with, or dissociate IGMP from, an interface.

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 deletes any saved parameters for that interface.

To display only those interfaces with IGMP associated, from the VLAN Interface Configuration mode enter the command `show igmp snooping/interface`.

No difference in behavior exists 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/SV L 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`. Because no VLANs are associated with IGMP, no output exists for this command.

To associate IGMP with VLAN 1, enable IGMP snooping, and then 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 MaxRsp T	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 two 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.

Configuring IGMP snooping on a VLAN

Enable IGMP snooping on a VLAN to forward the multicast data to only those ports that are members of the group.

IGMP snooping is disabled by default.

Procedure steps

1. Log on to VLAN Interface Configuration command mode in ACLI.
2. Enable IGMP snooping:

```
[default] [no] ip igmp snooping
```

OR

1. Log on to Global Configuration command mode in ACLI:
2. Enable IGMP snooping:

```
[default] vlan igmp <vid> [snooping {enable | disable}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Disables IGMP snooping on the selected VLAN.
no	Disables IGMP snooping on the selected VLAN.
enable	Enables IGMP snooping on the selected VLAN.
disable	Disables IGMP snooping on the selected VLAN.

Configuring IGMP send query on a VLAN

Use this procedure to enable IGMP send query on a snoop-enabled VLAN. If you enable IGMP snooping send query, 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 forwarding.

IGMP send query is disabled by default.

Prerequisites

- You must enable snoop on the VLAN.

Procedure steps

1. Log on to VLAN Interface Configuration mode in ACLI.
2. Enable IGMP send query:

```
ip igmp send-query
```

Configuring IGMP proxy on a VLAN

Use this procedure to enable IGMP proxy on a snoop-enabled VLAN. With IGMP proxy enabled, the switch consolidates incoming report messages into one proxy report for that group.

IGMP proxy is disabled by default.

Prerequisites

- You must enable snoop on the VLAN.

Procedure steps

1. Log on to VLAN Interface Configuration mode in ACLI.
2. Enable IGMP proxy:

```
[default] [no] ip igmp proxy
```

OR

1. Log on to Global Configuration command mode in ACLI.
2. Enter the following:

```
[default] [no] vlan igmp <vid> [proxy {enable | disable}]
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Disables IGMP proxy on the selected VLAN.
no	Disables IGMP proxy on the selected VLAN.
<vid>	Specifies the VLAN ID.
enable	Enables IGMP proxy on the selected VLAN.
disable	Disables IGMP proxy on the selected VLAN.

Configuring the IGMP version on a VLAN

Use this procedure to configure the IGMP version to run on the VLAN. You can specify the version as IGMPv1, IGMPv2, or IGMPv3. The default is IGMPv2.

Procedure steps

1. Log on to VLAN Interface Configuration mode in ACLI.
2. Configure the IGMP version:

```
[default] ip igmp version <1-3>
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Restores the default IGMP protocol version (IGMPv2).
<1-3>	Specifies the IGMP version.

Configuring static mrouter ports on a VLAN

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.

Procedure steps

1. Log on to VLAN Interface Configuration command mode in ACLI.
2. Configure static mrouter ports on a VLAN (IGMPv1, IGMPv2, and IGMPv3 according to the supported version on the VLAN):

```
[default] [no] ip igmp mrouter <portlist>
```

OR

1. Log on to Global Configuration command mode in ACLI.
2. Configure IGMPv1 or IGMPv2 static mrouter ports:

```
[no] vlan igmp <vid> {v1-members | v2-members} [add | remove]
<portlist>
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Removes all static mrouter ports.
no	Removes the specified static mrouter port.
<portlist>	Specifies the list of ports to add or remove as static mrouter ports.
{v1-members v2-members}	Specifies whether the static mrouter ports are IGMPv1 or IGMPv2.
[add remove]	Specifies whether to add or remove the static mrouter ports.
<portlist>	Specifies the list of ports to add or remove as static mrouter ports.

Displaying IGMP snoop, proxy, and mrouter configuration

Use this procedure to display the IGMP snoop, proxy, and mrouter configuration for each VLAN.

Procedure steps

To display IGMP snoop information, enter:

```
show ip igmp snooping
```

Job aid

The following table shows the field descriptions for the `show ip igmp snooping` command.

Field	Description
Vlan	Indicates the Vlan ID.
Snoop Enable	Indicates whether snoop is enabled (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 ages 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.

Configuring IGMP parameters on a VLAN

Use this procedure to configure the IGMP parameters on a VLAN.

! Important:

The query interval, robustness, and version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

Procedure steps

1. Log on to VLAN Interface Configuration command mode in ACLI.
2. Configure IGMP parameters:

```
[default] ip igmp
[last-member-query-interval <last-mbr-query-int>]
[query-interval <query-int>]
[query-max-response <query-max-resp>]
[robust-value <robust-val>]
[version <1-3>]
```

OR

1. Log on to Global Configuration command mode in ACLI.
2. Configure IGMP parameters:

```
[default] vlan igmp <vid>
[query-interval <query-int>]
[robust-value <robust-val>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Configures the selected parameter to the default value. If you do not specify a parameter, snoop is disabled and all IGMP parameters use their default values.
<last-mbr-query-int>	Configures the maximum response time (in 1/10 seconds) 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. You cannot configure this value 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 (1 second). 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. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)
<query-int>	Configures the frequency (in seconds) at which host query packets are transmitted on the VLAN. The range is 1–65535. The default value is 125 seconds.
<query-max-resp>	Specifies the maximum response time (in 1/10 seconds) advertised in IGMPv2 generalqueries on this interface. The range is 0–255. The default value is 100 (10 seconds).
<robust-val>	Specifies 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.

Configuring the router alert option on a VLAN

Use this command 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

Procedure steps

1. Log on to VLAN Interface Configuration mode in ACLI.
2. Configure the router alert option on a VLAN:

```
[default] [no] ip igmp router-alert
```

Variable definitions

The following table describes the command variables.

Variable	Value
default	Disables the router alert option.
no	Disables the router alert option.

Displaying IGMP interface information

Use this procedure to display IGMP interface parameters.

Procedure steps

To display the IGMP interface information, enter:

```
show ip igmp interface [vlan <vid>]
```

OR

Enter:

```
show vlan igmp <vid>
```

Job aids

The following table shows the field descriptions for the `show ip igmp interface` command.

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 attaches.
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. This value is also the amount of time between group-specific query messages.

Field	Description
	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.
Send Query	Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES or NO. The default is disabled.

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

Field	Description
Snooping	Indicates whether snooping is enabled or disabled.
Proxy	Indicates whether proxy snoop is enabled or disabled.
Robust Value	Indicates the robust value configured for expected packet loss on the interface.
Query Time	Indicates the frequency (in seconds) at which host query packets are transmitted on the interface.
IGMPv1 Static Router Ports	Indicates the IGMPv1 static mrouter ports.
IGMPv2 Static Router Ports	Indicates the IGMPv2 static mrouter ports.
Send Query	Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES or NO. The default is disabled.

Displaying IGMP Multicast filtering mode

Use this procedure to display the IGMP multicast filtering mode.

Procedure steps

1. Log on to the Privileged EXEC Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
show ip igmp multicast-filter-mode
```

Configuring IGMP multicast filtering mode

Use this procedure to configure the IGMP multicast filtering mode.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip igmp multicast-filter-mode
```

to enable IGMP multicast filtering mode.
OR

```
no ip igmp multicast-filter-mode
```

to disable.

*** Note:**

By default, this feature is disabled. You can reset this feature to the default from any state by entering the following command: `default ip igmp multicast-filter-mode`

Displaying IGMP group membership information

Display the IGMP group information to show the learned multicast groups and the attached ports.

Procedure steps

To display IGMP group information, enter:

```
show ip igmp group [count] [group <A.B.C.D>] [member-subnet <A.B.C.D>/<0-32>]
```

OR

Enter:

```
show vlan multicast membership <vid>
```

Variable definitions

The following table describes the command variables.

Variable	Value
count	Displays the number of IGMP group entries.

Variable	Value
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 aids

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 after receiving a group report.
Type	Specifies the type of membership: static or dynamic.
In Port	Identifies the member port for the group. This port is the port on which group traffic is forwarded and in cases where the type is dynamic, it is the port on which the IGMP join was received.

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.

Configuring unknown multicast packet filtering

The default switch behavior is to flood all packets with unknown multicast addresses. Use this procedure to prevent the flooding of packets with unknown multicast addresses, and to enable the forwarding of these packets to static mrouter ports only.

Procedure steps

1. Log on to Global Configuration mode in ACLI.
2. Configure unknown multicast packet flooding:

```
[no] [default] vlan igmp <vid> unknown-mcast-no-flood {enable  
| disable}
```

Variable definitions

The following table describes the command variables.

Variable	Value
no	Enables the flooding of multicast packets on the VLAN.
default	Enables the flooding of multicast packets on the VLAN.
enable	Prevents the flooding of multicast packets on the VLAN.
disable	Enables the flooding of multicast packets on the VLAN.

Displaying the status of unknown multicast packet filtering

Use this procedure to display the status of unknown multicast filtering: enabled (no flooding) or disabled (flooding allowed).

Procedure steps

To display the unknown multicast flooding configuration, enter:

```
show vlan igmp unknown-mcast-no-flood
```

Job aid

The following table shows the field descriptions for the `show vlan igmp unknown-mcast-no-flood` command.

Field	Description
Unknown Multicast No-Flood	Specifies the status of unknown multicast filtering: enabled or disabled.

Specifying a multicast MAC address to be allowed to flood a VLAN

Use this procedure to allow particular unknown multicast packets to be flooded on a VLAN.

To add MAC addresses that start with 01.00.5E to the allow-flood table, you must specify the corresponding multicast IP address. For instance, you cannot add MAC address 01.00.5E.01.02.03 to the allow-flood table; instead you must specify IP address 224.1.2.3.

For all other types of MAC address, you can enter the MAC address directly to allow flooding.

Procedure steps

1. Log on to Global Configuration mode in ACLI.
2. Allow particular unknown multicast packets to be flooded:

```

vlan igmp unknown-mcast-allow-flood <vid> {<H.H.H> |
<mcast_ip_address>}

```

Variable definitions

The following table describes the command variables.

Variable	Value
<vid>	Specifies the VLAN on which to configure the flooding.
<H.H.H>	Specifies the multicast MAC address to flood. Acceptable formats are: <ul style="list-style-type: none"> • H.H.H • xx:xx:xx:xx:xx:xx • xx.xx.xx.xx.xx.xx • xx-xx-xx-xx-xx-xx
<mcast_ip_address>	Specifies the multicast IP address to flood.

Displaying the multicast MAC addresses for which flooding is allowed

Use this procedure to display the multicast MAC addresses for which flooding is allowed on all switch VLANs.

Procedure steps

To display the multicast MAC addresses for which flooding is allowed, enter:

```
show vlan igmp unknown-mcast-allow-flood
```

Job aid

The following table shows the field descriptions for the `show vlan igmp unknown-mcast-allow-flood` command.

Field	Description
Allowed Multicast Addresses	Indicates multicast addresses that can flood.

Displaying IGMP cache information

Display the IGMP cache information to show the learned multicast groups in the cache and the IGMPv1 version timers.

*** Note:**

Using the `show ip igmp cache` command does not display the expected results in some configurations. If you do not see the expected results, use the `show ip igmp group` command to view the information.

Procedure steps

To display the IGMP cache information, enter:

```
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. After the interface hears 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.

Flushing the router table

Use this procedure to flush the router table.

Procedure steps

1. Log on to Global Configuration mode in ACLI.
2. Flush the router table:

```
ip igmp flush vlan <vid> {grp-member|mrouter}
```

Variable definitions

The following table describes the command variables.

Variable	Value
{grp-member mrouter}	Flushes the table specified by type.

Configuring IGMP selective channel block

In certain deployment scenarios, you may need to prevent multicast streaming from specific group addresses to users that connect to certain ports. You can use the IGMP selective channel block feature to prevent this streaming. IGMP selective channel block controls the IGMP membership of ports by blocking IGMP reports received from users on that port and destined for the specific group address or addresses. You can configure the filter to block a single multicast address or a range of addresses.

This feature works regardless of whether the switch is in Layer 2 IGMP snooping mode or the full IGMP mode (PIM-SM enabled). This feature also applies to IGMPv1 and v2.

Creating an IGMP profile

Use this procedure to create an IGMP profile.

1. Log on to Global Configuration mode in ACLI.
2. Enter the following command:


```
ip igmp profile <profile number (1-65535)>
```
3. Enter the following command:


```
range <ip multicast address> <ip multicast address>
```

Deleting an IGMP profile

Use this procedure to delete an IGMP profile.

1. Log on to Global Configuration mode in ACLI.
2. Delete an IGMP profile:


```
no ip igmp profile <profile number (1-65535)>
```

Applying the IGMP filter profile on interface

Use this procedure to apply the IGMP filter profile on an interface.

1. Log on to Global Configuration mode in ACLI.
2. Enter the following command:

```
interface <interface-id>
```
3. Enter the following command:

```
ip igmp filter <profile number>
```

Removing a profile from an interface

Use this procedure to remove a profile from an interface.

1. Log on to Global Configuration mode in ACLI.
2. Enter the following command:

```
interface <interface-id>
```
3. Enter the following command:

```
no ip igmp filter <profile number>
```

Displaying an IGMP profile

Use this procedure to display an IGMP profile.

1. Log on to Global Configuration mode in ACLI.
2. Display an IGMP profile:

```
show ip igmp profile <cr> or <profile number>
```


Chapter 36: IGMP snooping configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure IGMP snooping on a VLAN using Enterprise Device Manager (EDM).

Configuring IGMP snooping using EDM

Enable snoop on the VLAN.

All related configurations, listed below, are optional and can be configured to suit the requirements of your network.

Configuring IGMP snoop, proxy, and IGMP parameters on a VLAN using EDM

Use the following procedure to configure IGMP snoop, proxy, and IGMP parameters on a VLAN.

IGMP snoop and proxy are disabled by default.

! **Important:**

With IGMP snoop, the QueryInterval, Robustness, and Version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

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 **Snoop** tab.

4. In the table, double-click the cell under the **Enable** column heading.
5. Select true from the drop-down list to enable IGMP snoop for a VLAN.
6. In the table, double-click the cell under the **ReportProxyEnable** column heading.
7. Select true from the drop-down list to enable IGMP proxy for a VLAN.
8. Repeat these steps for the other column headings until you have amended all of the parameters you want to change.
9. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Snoop** tab.

Field	Description
Id	Specifies the VLAN ID.
Name	Specifies the VLAN name.
ReportProxyEnable	Specifies the IGMP proxy status: enabled (true) or disabled (false).
Enable	Specifies the IGMP snoop status: enabled (true) or disabled (false).
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each 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 2–255, 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.
QueryInterval	Specifies the frequency (in seconds) at which IGMP query packets are transmitted on the VLAN. Ensure that the query interval is the same as the configured value on the multicast router (IGMP querier). The range is from 1 to 65535, and the default is 125.
MRouterPorts	Specifies the statically-configured mrouter ports in the VLAN that provide connectivity to a nonquerier IP multicast router. Multicast data and group reports are forwarded out this port to the multicast router. You do not need to configure this parameter if only one multicast router exists on the VLAN.
Ver1MRouterPorts	Displays the mrouter ports in the VLAN that use IGMP version 1.

Field	Description
Ver2MRouterPorts	Displays the mrouter ports in the VLAN that use IGMP version 2.
ActiveMrouterPorts	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
ActiveQuerier	Displays the IP address of the multicast querier router.
QuerierPort	Displays the mrouter port on which the multicast querier router was heard.
MrouterExpiration	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.

Configuring IGMP snoop, proxy, and static mrouter ports on a VLAN using EDM

Use the following procedure to configure IGMP snooping, proxy, and static mrouter ports on a VLAN.

By default, IGMP snoop and proxy are disabled, and no static mrouter ports are configured.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IGMP**.
3. In the work area, click the **Snoop** tab.
4. In the table, double-click the cell under the **SnoopEnable** column heading.
5. Select true from the drop-down list to enable IGMP snoop.
6. In the table, double-click the cell under the **ProxySnoopEnable** column heading.
7. Select true from the drop-down list to enable IGMP proxy.
8. In the table, double-click the cell under the **SnoopMRouterPorts** column heading.
9. Select the desired ports from the list to configure mrouter ports.
10. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Snoop** tab.

Field	Description
IfIndex	Specifies the VLAN ID.
SnoopEnable	Specifies the IGMP snoop status: enabled (true) or disabled (false).
ProxySnoopEnable	Specifies the IGMP proxy status: enabled (true) or disabled (false).
SnoopMRouterPorts	Specifies 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.
SnoopActiveMRouterPorts	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
SnoopMRouterExpiration	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.

Flushing the IGMP router tables and configuring IGMP router alert using EDM

Use the following procedure to flush the IGMP router tables. This procedure also allows you to configure the router alert option and query interval on a VLAN.

! **Important:**

The QueryInterval, Robustness, and Version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

! Important:

To maximize your network performance, Avaya recommends that you set the router alert parameter according to the version of IGMP currently in use:

- IGMPv1 - Disable
- IGMPv2 - Enable
- IGMPv3 - Enable

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IGMP**.
3. In the work area, click the **Interface** tab.
4. In the table, double-click the cell under the **FlushAction** column heading.
5. Select the desired flush option to flush the routing table.
6. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Interface** tab.

Field	Description
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.

Field	Description
QueryMaxResponseTime	Indicates the maximum response time (in 1/10 seconds) advertised in IGMPv2 general queries on this interface. This field is a read-only field and is not configurable on the Ethernet Routing Switch 5000 Series .
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 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	Specifies 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.
LastMemberQueryInterval	Sets 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	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 to send query or not.
FlushAction	Flushes the specified table type: <ul style="list-style-type: none"> • none • flushGrpMem: group member table • flushMrouter: mrouter table

Configuring unknown multicast filtering using EDM

The default switch behavior is to flood all packets with unknown multicast addresses.

Use the following procedure to prevent the flooding of packets with unknown multicast addresses and enable the forwarding of these packets to static mrouter ports only.

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 **Unknown Multicast Filtering** tab.
4. Select the **UnknownMulticastNoFlood** check box to enable unknown multicast flooding.
5. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Unknown Multicast Filtering** tab.

Field	Definition
Id	Indicates the VLAN ID.
UnknownMulticastNoFlood	Enables or disables unknown multicast flooding.

Specifying a multicast MAC address to be allowed to flood all VLANs using EDM

Use the following procedure to allow particular unknown multicast packets to be flooded on all switch 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 **MAC Multicast Filter Table** tab.
4. In the toolbar, click **Insert**.
5. Type the VLAN ID to flood in the **AllowedAddressVLANId** field.
6. Type the MAC address to flood in the **AllowedAddressMacAddr** field.
7. Click **Insert**.

Variable definitions

The following table describes the fields of the **MAC Multicast Filter Table** tab.

Field	Definition
AllowedAddressVLANId	Indicates the allowed VLAN ID.
AllowedAddressMacAddr	Indicates the allowed MAC address.

Specifying an IP address to be allowed to flood a VLAN using EDM

Use this procedure to configure the IP address multicast filter table. This table specifies multicast IP addresses that are allowed to be flooded to all ports on a per-VLAN basis.

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 **IP Address Multicast Filter Table** tab.
4. Click **Insert**

5. Complete the fields as required.
6. Click **Insert**

Variable definitions

The following table describes the fields of the **IP Address Multicast Filter Table** tab.

Field	Definition
VlanAllowedNetAddressVlanId	Specifies the ID of the VLAN to configure.
VlanAllowedNetAddressType	Specifies the address type: ipv4.
VlanAllowedNetAddress	Specifies a multicast IP address that is allowed to flood all ports. Unicast and broadcast addresses are not allowed.

Configuring SSM for IGMP

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.

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

Displaying the SSM mapping table

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.

Variable	Value
Activity	Displays SSM map activity.
AdminState	Indicates whether SSM mapping is enabled or disabled.

Creating an SSM map for IGMP

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

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 the SSM map source.
8. On the menu bar, click **Apply**.

Displaying IGMP cache information using EDM

Use the following procedure to display IGMP cache information to show the learned multicast groups in the cache and the IGMPv1 version timers.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IGMP**.
3. In the work area, click the **Cache** tab to view the IGMP cache information.

Variable definitions

The following table describes the fields of the **Cache** tab.

Field	Description
Address	Indicates the IP multicast group address.
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.

Field	Description
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.

Displaying IGMP group information using EDM

Use the following procedure to display IGMP group information to show the learned multicast groups and the attached ports.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IGMP**.
3. In the work area, click the **Groups** tab.

Variable definitions

The following table describes the fields of the **Groups** tab.

Field	Description
IpAddress	Indicates the multicast group 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.

Field	Description
InPort	Indicates the member port for the group. This is the port on which group traffic is forwarded.

Displaying extended interface IGMP group information

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 multicast route information using EDM

Use the following procedure to display multicast route information for troubleshooting purposes.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **Multicast**.
3. In the work area, click the **Routes** tab to view multicast routes information.

Variable definitions

The following table describes the fields of the **Routes** tab.

Field	Description
Group	Indicates the IP Multicast group address.
Source	Indicates the source address.
SourceMask	Indicates the source address mask.
UpstreamNeighbor	Indicates the address of the upstream neighbor that is forwarding packets for the specified source and group. 0.0.0.0 appears if the network is local.
Interface	Indicates the VLAN where datagrams for the specified source and group are received.
ExpiryTime	Indicates the amount of time remaining before this entry is aged out. The value 0 indicates that the entry is not subject to aging.
Protocol	Indicates the routing protocol through which this route was learned.

Displaying multicast next hop information using EDM

Use the following procedure to display all multicast next hop information to find the best route to a member group.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **Multicast**.
3. In the work area, click the **Next Hops** tab to view multicast next hops information.

Variable definitions

The following table describes the fields of the **Next Hops** tab.

Field	Description
Group	Indicates the IP Multicast group.
Source	Indicates the source address.
SourceMask	Indicates the source address mask.
OutInterface	Indicates the VLAN ID for the outgoing interface for the next hop.
Address	Indicates the address of the next hop specific to this entry. For most interfaces, this address is identical to the next hop group.
State	Indicates whether or not the outgoing interface and next hop represented by this entry is currently being used to forward IP datagrams. A Value of <i>forwarding</i> indicates this parameter is currently being used; <i>pruned</i> indicates the parameter is not being used.
ExpiryTime	Indicates the amount of time remaining before this entry is aged out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	Indicates the minimum number of hops between this router and a member of this IP Multicast group reached through this next hop on this outgoing interface. IP Multicast datagrams for the group that have a TTL less than this number of hops are not forwarded to the next hop.
Protocol	Indicates the routing protocol where this next hop was learned.

Displaying multicast interface information using EDM

Use the following procedure to display multicast interface information.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **Multicast**.
3. In the work area, click the **Interfaces** tab to view multicast interfaces information.

Variable definitions

The following table describes the fields of the **Interfaces** tab.

Field	Description
Interface	Indicates the VLAN ID.
Ttl	Indicates the datagram time-to-live (TTL) threshold for the interface. IP Multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 1 means that all multicast packets are forwarded out of the interface.
Protocol	Indicates the routing protocol running on this interface.

Creating an IGMP profile

Create an IGMP profile to configure the IGMP selective channel block feature.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IGMP**.
3. In the work area, click the **Profile** tab.
4. In the toolbar, click **Insert**.

The Insert Profile dialog box appears.

5. Type the profile ID in the **Profile ID** field.
6. Click **Insert**.

The table is updated with the created profile.

7. Double-click the **ProfilePortList** field to assign ports to the profile.

Variable definitions

The following table describes the fields of the **IGMP Profile Range** tab.

Field	Definition
ProfileId	Indicates the Profile ID. The range is from 1 to 65535.
ProfileType	Indicates the type of the profile.
ProfilePortList	Specifies the list of ports to which this profile applies.
ProfileDroppedPackets	Indicates the number of packets that were matched by this profile and dropped.

Configuring the IGMP Profile Range

Use the following procedure to configure the IGMP profile range.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **IGMP**.
3. In the work area, click the **Profile** tab.
4. Select an existing profile.
5. In the toolbar, click **Profile Range**.

The **Profile-Profile Range** tab appears.

6. In the table, double-click the cell under the column heading for the parameter you want to change.
7. Select a parameter or value from the drop-down list.

8. Repeat the previous two steps until you have amended all of the parameters you want to change.
9. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **IGMP Profile Range** tab.

Field	Definition
ProfileId	Indicates the Profile ID. The range is from 1 to 65535.
RangeAddressStart	Indicates a valid multicast group address as the start of the range.
RangeAddressEnd	Indicates a valid multicast group address as the end of the range.

Chapter 37: PIM configuration using ACLI

This chapter describes the procedures you can use to configure PIM-SM using ACLI.

Unlike dense-mode protocols, such as DVMRP that initially flood multicast traffic to all routers over an entire internetwork, PIM sends multicast traffic only to routers that belong to a specific multicast group and that choose to receive the traffic. PIM reduces overhead costs for processing unwanted multicast packets.

Prerequisites for PIM configuration

Before you can configure PIM, you must prepare the switch as follows:

1. Install the Advanced Routing software license.

! Important:

If your Ethernet Routing Switch is running an Advanced License for a release prior to Release 6.0, to enable PIM-SM you must regenerate your license file from the Avaya web site and install the new license file on the switch.

2. Enable routing globally.
3. Configure IP addresses and enable routing on the VLAN interfaces on which you want to configure PIM-SM.
4. Enable a unicast protocol, either RIP or OSPF, globally and on the interfaces on which you want to configure PIM.

! Important:

PIM requires a unicast protocol to multicast traffic within the network when performing the Reverse Path Forwarding (RPF) check. PIM also uses the information from the unicast routing table to create and maintain the shared and shortest path multicast tree. The unicast routing table must contain a route to every multicast source in the network, as well as routes to PIM entities such as the rendezvous points (RP) and bootstrap router (BSR).

PIM configuration procedures

To configure PIM-SM, you must perform the following procedures:

1. Enable PIM globally.
(If desired, modify the default global PIM properties.)
2. Enable PIM on individual VLAN interfaces.
(If desired, modify the default VLAN PIM properties.)
3. For PIM-SM, configure candidate RPs for the multicast groups in the network. (It is best to have multiple candidate-RPs in the network; however, with the Ethernet Routing Switch 5000, you can only configure one candidate-RP per switch for any number of groups.)

OR

Configure one (or several) static RPs for the multicast groups in the network. (To enable static RP in the PIM-SM domain, you must configure the same static RPs on every system that takes part in PIM-SM forwarding.)

4. For PIM-SM, configure one or several candidate BSRs to propagate RP information to all switches in the network. (You can configure every PIM-enabled VLAN as a C-BSR. If Static RP is enabled, this step is not required.)

! Important:

Ensure that all routers in the path from the receivers to the RP and to the multicast source are PIM-enabled. Also ensure that all PIM routers have unicast routes to reach the source and RP through directly-connected PIM neighbors.

Required configuration steps

To configure PIM-SSM, you must perform the following procedures:

1. Enable PIM globally and change PIM mode to SSM.
(If desired, modify the default global PIM properties.)
2. Enable PIM on individual VLAN interfaces.
(If desired, modify the default VLAN PIM properties.)
3. If you use PIM-SSM with the IGMPv3 protocol, then configure this option on each VLAN.

All additional configurations listed below are optional and can be configured according to the requirements of your network.

Job aid: Roadmap of PIM configuration commands

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

Command	Parameter
Global Configuration mode	
ip pim	bootstrap-period <bootstrap-period >
	rp-c-adv-timeout <rp-c-adv-time>
	disc-data-timeout <disc-data-time>
	enable
	mode <pim-mode>
	join-prune-interval <join-prune-int>
	fwd-cache-timeout <fwd-cache-time>
	register-suppression-timeout <rgstr-suppr-time>
	unicast-route-change-timeout <unicast-rte-chge-time>
ip pim rp-candidate	group <group-addr> <group-mask> rp <rp-addr>
ip pim static-rp	enable
	<group-addr> <group-mask> <static-rp-addr>
ip pim virtual-neighbor	<if-ipaddr> <v-nbr-ipaddr>
Interface vlan mode	
ip pim	bsr-candidate priority <priority>
	enable
	interface-type <active passive>
	join-prune-interval <join-prune-int>
	query-interval <query-int>
Privileged EXEC mode	
show ip mroute	{interface next-hop route}
show ip pim	
show ip pim active-rp	[group <group-addr>]

Command	Parameter
show ip pim bsr	
show ip pim interface	[vlan <vlan-id>]
show ip pim mode	
show ip pim mroute	[source <ipaddr>] [group <group>] [summary]
show ip pim neighbor	
show ip pim rp-candidate	[group <group-addr>]
show ip pim rp-hash	
show ip pim static-rp	
show ip pim virtual-neighbor	

Enabling and disabling PIM-SM globally

Use this procedure to enable PIM-SM on individual interfaces, you must first enable PIM-SM globally.

By default, PIM-SM is disabled.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip pim enable
```

Variable definitions

The following table describes the command variable.

Variable	Value
[no]	Disables PIM-SM globally.

Enabling and disabling PIM-SSM globally

Use this procedure to enable or disable PIM-SSM. To enable PIM-SSM on individual interfaces, you must first enable PIM-SSM globally. By default PIM-SSM is disabled.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command to enable PIM-SSM:

```
ip pim enable mode ssm
```
3. At the command prompt, enter the following command to disable PIM-SSM:

```
no ip pim [enable]
```

Configuring global PIM-SM properties

Use this procedure to configure the global PIM-SM parameters on the switch.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command to configure the PIM bootstrap period:

```
ip pim bootstrap-period <bootstrap-period>
```
3. At the command prompt, enter the following command to configure the PIM discard data timeout:

```
ip pim disc-data-timeout <disc-data-time>
```
4. At the command prompt, enter the following command to configure the forward cache timeout globally:

```
ip pim fwd-cache-timeout <fwd-cache-time>
```
5. At the command prompt, enter the following command to configure the join-prune interval:

```
ip pim join-prune-interval <join-prune-int>
```

6. At the command prompt, enter the following command to configure the PIM mode globally:

```
ip pim mode <pim-mode>
```

7. At the command prompt, enter the following command to configure the register suppression timeout:

```
ip pim register-suppression-timeout <rgstr-suppr-time>
```

8. At the command prompt, enter the following command to configure the how often the candidate RPs send C-RP advertisement messages:

```
ip pim rp-c-adv-timeout <rp-c-adv-time>
```

9. At the command prompt, enter the following command to configure the PIM-SM unicast route change timeout:

```
ip pim unicast-route-change-timeout <unicast-rte-chge-time>
```

Variable definitions

The following table describes the `ip pim` command variables.

Variable	Value
bootstrap-period <bootstrap-period>	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. Range is 5–32757. The default is 60.
disc-data-timeout <disc-data-time>	After the router forwards the first source packet to the RP, this value specifies how long (in seconds) the router discards subsequent source data while waiting for a join from the RP. An IPMC discard record is created and deleted after the timer expires or after a join is received. Range is 5–65535. The default is 60.
fwd-cache-timeout <fwd-cache-time>	Specifies the forward cache timeout globally. This value is used in aging PIM-SM mroutes. Range is 10–86400. The default is 210.
join-prune-interval <join-prune-int>	Specifies how long to wait (in seconds) before the PIM-SM router sends out the next join/prune message to the upstream neighbors. Range is 1–18724. The default is 60.
mode <pim-mode>	Specifies sparse or ssm mode.
register-suppression-timeout <rgstr-suppr-time>	Specifies the PIM-SM register suppression timeout. Range is 6–65535. The default is 60.

Variable	Value
<code>rp-c-adv-timeout <rp-c-adv-time></code>	Specifies how often (in seconds) candidate RPs (C-RP) send C-RP advertisement messages. After this timer expires, the C-RP sends an advertisement message to the elected BSR. Range is 5–26214. The default is 60.
<code>unicast-route-change-timeout <unicast-rte-chge-time></code>	Specifies the PIM-SM unicast route change timeout. Indicates how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates to be used by PIM. Range is 2–65535. The default is 5.

Displaying global PIM-SM properties

Use this procedure to display global PIM-SM properties.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim
```

Job aid

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

Field	Description
PIM Admin Status	Indicates the status of PIM-SM.
PIM Boot Strap Period	Indicates the interval between originating bootstrap messages at the elected BSR.
PIM C-RP-Adv Message Send Interval	Indicates the candidate RPs timer (in seconds) for sending C-RP advertisement messages.
PIM Discard Data Timeout	After the router forwards the first source packet to the RP, this value indicates how long (in seconds) the router discards subsequent source data while waiting for a join from the RP. An IPMC discard record is created and deleted after the timer expires or after a join is received.
PIM Join Prune Interval	Indicates the join/prune interval in seconds.

Field	Description
PIM Register Suppression Timer	Indicates the register suppression timer in seconds.
PIM Uni Route Change Timeout	Indicates how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates to be used by PIM.
PIM Mode	Indicates the PIM mode: sparse mode or source specific multicast mode.
PIM Static-RP	Indicates the status of static RP.
Forward Cache Timeout	Indicates the PIM-SM forward cache expiry value in seconds. This value is used in aging PIM-SM mroutes.

Enabling PIM-SM on a VLAN

Use this procedure to enable PIM-SM on a VLAN.

By default, PIM-SM is disabled on VLANs.

Prerequisites

- You must enable PIM-SM globally.

Procedure steps

1. Enter VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip pim enable
```

Variable definitions

The following table describes the command variable.

Variable	Value
[no]	Disables PIM-SM on the VLAN.

Configuring the PIM-SM interface type on a VLAN

Use this procedure to change the state (active or passive) of PIM on a VLAN interface. An active interface transmits and receives PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you have a high number of PIM interfaces and these interfaces are connected to end users, not to other switches.

By default, VLANs are active interfaces.

Prerequisites

- Before you change the state of PIM on a VLAN interface, you must disable PIM on the interface to prevent instability in the PIM operations, especially when neighbors are present or when streams are received.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip pim interface-type <active|passive>
```

Variable definitions

The following table describes the command variables.

Variable	Value
interface-type <active passive>	Sets the interface type on a particular VLAN: <ul style="list-style-type: none"> • active: allows PIM-SM control traffic to be transmitted and received. • passive: prevents PIM-SM control traffic from being transmitted or received, reducing the load on the system.

Displaying PIM-SM neighbors

Use this procedure to display PIM-SM neighbors.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim neighbor
```

Job aid

The following table shows the field descriptions for the `show ip pim neighbor` command.

Field	Description
Address	Specifies the IP address of the PIM-SM neighbor.
Vlan	Specifies the local interface.
Uptime	Specifies the elapsed time since the PIM-SM neighbor last became a neighbor of the local interface.
Expiry Time	Specifies the time remaining before this PIM-SM neighbor times out.
Total PIM Neighbors	Specifies the total number of PIM neighbors on the switch.

Configuring PIM-SM properties on a VLAN

Use this procedure to configure the PIM-SM properties on a VLAN to modify the join/prune interval or the query interval.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip pim join-prune-interval <join-prune-int> query-interval
<query-int>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<join-prune-int>	Specifies how long to wait (in seconds) before the PIM-SM switch sends out the next join/prune message to the upstream neighbors. Range is 1–18724, and the default is 60.
<query-int>	Sets the hello interval for the VLAN. The range is 0–18724. The default is 30.

Displaying the PIM-SM configuration for a VLAN

Use this procedure to display information about the PIM-SM interface configuration for a VLAN.

Procedure steps

1. Log on to PIM-SM Interface Configuration mode.
2. At the command prompt, enter the following command:

```
show ip pim interface [vlan <vid>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
<vid>	Specifies the VLAN to display (1–4094).

Job aid

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

Field	Description
Vlan	Identifies the VLAN.
State	Indicates the state of PIM-SM on the VLAN.
Address	Specifies the VLAN IP address.
Mask	Specifies the VLAN subnet mask.
Mode	Indicates the PIM mode of this VLAN: sparse mode or source specific multicast mode.
DR	Indicates the Designated Router for this interface.
Hello Interval	Indicates how long the switch waits (in seconds) between sending out a hello message to neighboring switches. The default hello interval is 30 seconds.
Join Prune Interval	Indicates how long the switch waits (in seconds) between sending out a join/prune message to the upstream neighbors. The default join/prune interval is 60 seconds.
CBSPR	Indicates the priority for this local interface to become a Candidate BSR. The Candidate BSR with the highest BSR priority and address is referred to as the preferred BSR. The default is -1, which indicates that the current interface is not a Candidate BSR.
Oper State	Indicates the status of PIM-SM on this interface: up or down.
Interface Type	Indicates whether the PIM-SM interface is active or passive.

Specifying the router as a candidate BSR on a VLAN

Because PIM-SM cannot run without a bootstrap router (BSR), you must specify at least one C-BSR in the domain. The C-BSR with the highest configured priority becomes the BSR for the domain. You can configure additional C-BSRs to provide backup protection in case the primary BSR fails.

If two C-BSRs have equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with the highest priority to the domain, it automatically becomes the new BSR.

With the Ethernet Routing Switch 5000 Series, you can configure every PIM-enabled interface as a C-BSR.

Procedure steps

1. Log on to the VLAN Interface Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip pim bsr-candidate priority <priority>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<priority>	Specifies the priority value of the candidate to become a BSR. The range is 0 to 255 and the default is -1, which indicates that the current interface is not a Candidate BSR.
[no]	Removes the candidate BSR configuration.

Displaying the BSR configuration

Use this procedure to display the current BSR configuration.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim bsr
```

Job aid

The following table shows the field descriptions for the `show ip pim bsr` command.

Field	Description
Current BSR Address	Specifies the IP address of the current BSR for the local PIM-SM domain.
Current BSR Priority	Specifies the priority of the current BSR. The Candidate BSR (C-BSR) with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
Current BSR Hash Mask	Specifies the mask used in the hash function to map a group to one of the C-RPs from the RP-Set. With the hash-mask, a small number of consecutive groups (for example, four) can always hash to the same RP.
Current BSR Fragment Tag	Specifies a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same Fragment Tag.
Current BSR Boot Strap Timer	Specifies the time the BSR waits between sending bootstrap messages.

Specifying a local IP interface as a candidate RP

Because PIM-SM cannot run without an RP, you must specify at least one C-RP in the domain. Use this procedure to configure a local PIM-SM interface as a candidate RP (C-RP).

With the Ethernet Routing Switch 5000 Series, you can configure only one local interface as a C-RP for any number of groups.

With the mask value, you can configure a C-RP for several groups in one configuration. For example, with a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0, you can configure the C-RP for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure steps

1. Log on to the Global Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip pim rp-candidate group <group-addr> <group-mask> rp
<c-rp-addr>
```

Variable definitions

The following table describes the command variables.

Variable	value
<group-addr>	Specifies the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
<group-mask>	Specifies the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
<c-rp-addr>	Specifies the IP address of the C-RP. This address must be one of the local PIM-SM enabled interfaces.
[no]	Removes the configured RP candidate.

Displaying the candidate RP configuration

Use this procedure to display the candidate RP configuration.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim rp-candidate [group <group-addr>]
```

Variable definitions

The following table describes the command variables.

Variable	value
<group-addr>	Specifies the IP address of the multicast group configuration to display.

Job aid

The following table shows the field descriptions for the `show ip pim rp-candidate` command.

Field	Description
Group Address	Specifies the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
Group Mask	Specifies the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RP Address	Specifies the IP address of the C-RP.

Displaying the PIM-SM RP set

Display the RP set for troubleshooting purposes. The BSR constructs the RP set from C-RP advertisements, and then distributes it to all PIM routers in the PIM domain for the BSR.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim rp-hash
```

Job aid

The following table shows the field descriptions for the `show ip pim rp-hash` command.

Field	Description
Group Address	Specifies the IP address of the multicast group.
Group Mask	Specifies the address mask of the multicast group.
Address	Specifies the IP address of the C-RP for the specified group.
Hold Time	Indicates the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for

Field	Description
	the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
Expiry Time	Specifies the time remaining before this C-RP times out.

Displaying the active RP per group

Use this procedure to display the active RP per group.

The active RP is displayed only when there is at least one (*,G) or (S,G) entry on the router after either joins or multicast data are received by the router.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim active-rp [group <group-addr>]
```

Variable definitions

The following table describes the command variables.

Variable	value
<group-addr>	Specifies the IP address of the multicast group configuration to display.

Job aid

The following table shows the field descriptions for the `show ip pim active-rp` command.

Field	Description
Group Address	Specifies the IP address of the multicast group.
Group Mask	Specifies the address mask of the multicast group.
Active RP	Specifies the IP address of the active RP.
Priority	Specifies the RP priority.

Enabling and disabling static RP

Enable static RP to avoid the process of dynamically learning C-RPs through the BSR mechanism. With this feature, static RP-enabled Ethernet Routing Switch 5000 Series switches can communicate with switches from other vendors that do not use the BSR mechanism.

! Important:

When you enable static RP, all dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure steps

1. Log on to the Global Configuration Mode in ACLI.
2. At the command prompt, enter the following command:

```
[no] ip pim static-rp [enable]
```
3. After you enter the command, a warning message appears. To confirm the change, enter:

```
y
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables static RP.
[enable]	Enables static RP.

Configuring a static RP

Use this procedure to configure a static RP entry. After you configure static RP, the switch ignores the BSR mechanism and uses only the RPs that you configure statically.

! Important:

You cannot configure a static RP-enabled switch as a BSR or as a C-RP.

Prerequisites

- You must enable static RP.

Procedure steps

At the command prompt, enter the following command:

```
[no] ip pim static-rp <group-addr> <group-mask> <static-rp-addr>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<group-addr>	Specifies the IP address of the multicast group. Together with the group mask, the group address identifies the range of the multicast addresses that the RP handles.
<group-mask>	Specifies the address mask of the multicast group. Together with the group address, the address mask identifies the range of the multicast addresses that the RP handles.
<static-rp-addr>	Specifies the IP address of the static RP.

Displaying the static RP configuration

Use this procedure to display the static RP configuration.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim static-rp
```

Job aid

The following table shows the field descriptions for the `show ip pim static-rp` command.

Field	Description
Group Address	Indicates the IP address of the multicast group. When combined with the group mask, the group address identifies the prefix that the local router uses to advertise as a static RP.
Group Mask	Indicates the address mask of the multicast group. When combined with the group address, the group mask identifies the prefix that the local router uses to advertise as a static RP.
RP Address	Indicates the IP address of the static RP.
Status	Indicates the status of static RP.

Specifying a virtual neighbor on an interface

Configure a virtual neighbor when the next hop for a static route cannot run PIM-SM, such as a Virtual Redundancy Router Protocol address on an adjacent device. The virtual neighbor IP address appears in the Ethernet Routing Switch 5000 neighbor table.

Procedure steps

At the command prompt, enter the following command:

```
[no] ip pim virtual-neighbor <if-ipaddr> <v-nbr-ipaddr>
```

Variable definitions

The following table describes the command variables.

Variable	Value
<if-ipaddr>	Specifies the IP address of the selected interface.
<v-nbr-ipaddr>	Specifies the IP address of the virtual neighbor.

Variable	Value
[no]	Removes the configured virtual neighbor.

Displaying the virtual neighbor configuration

Use this procedure to display the virtual neighbor.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim virtual-neighbor
```

Job aid

The following table shows the field descriptions for the `show ip pim virtual-neighbor` command.

Field	Description
Vlan	Indicates the VLAN interface.
Neighbor address	Indicates the IP address of the virtual neighbor.

Displaying the PIM mode

Use this procedure to display the PIM mode.

Procedure steps

At the command prompt, enter the following command:

```
show ip pim mode
```

Displaying multicast route information

Use this procedure to display multicast route information.

Procedure steps

At the command prompt, enter the following command:

```
show ip mroute {interface | next-hop | route}
```

Variable definitions

The following table describes the command variables.

Variable	Value
interface	Displays multicast route information per interface.
next-hop	Displays next-hop information for the multicast routes
route	Displays multicast route information.

Job aids

The following table shows the field descriptions for the `show ip mroute interface` command.

Field	Description
Interface	Indicates the interface.
Ttl	Indicates the datagram TTL threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means all multicast packets are forwarded out of the interface.
Protocol	Indicates the routing protocol running on this interface.

The following table shows the field descriptions for the `show ip mroute next-hop` command.

Field	Description
Interface	Indicates the interface identity.
Group	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
Source	Indicates the network address, which when combined with the corresponding value of Srcmask identifies the sources for which this entry specifies a next hop on an outgoing interface.
Srcmask	Indicates the network mask, which when combined with the corresponding value of Source identifies the sources for which this entry specifies a next hop on an outgoing interface.
Address	Indicates the address of the next hop specific to this entry. For most interfaces, this address is identical to Group.
State	Indicates whether the outgoing interface and next hop represented by this entry are currently forwarding IP datagrams. The value forwarding indicates the information is currently used. The value pruned indicates it is not used.
Exptime	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Closehop	Indicates the minimum number of hops between this router and members of this IP multicast group reached through this next hop on this outgoing interface. IP multicast datagrams for the group that use a TTL less than this number of hops are forwarded to the next hop
Protocol	Indicates the routing mechanism through which this next hop was learned.

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

Field	Description
Group	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
Source	Indicates the network address that, when combined with the corresponding value of Srcmask, identifies the sources for which this entry specifies a next hop on an outgoing interface.
Srcmask	Indicates the network mask that, when combined with the corresponding value of Source, identifies the sources for which this entry specifies a next hop on an outgoing interface.

Field	Description
Upstream_nbr	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is unknown.
If	Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).
Expir	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Prot	Indicates the outgoing mechanism through which this route was learned.

Chapter 38: PIM-SM/SSM configuration example using ACLI

Example 1

In this first example, A1 is an 8-unit mixed stack of ERS 5500 and ERS 5600 Series switches running IGMPv2 snooping.

A2, A3, and CW1 are all ERS 5600 Series switches with PIM-SM enabled.

RIP is used as the Layer 3 routing protocol but you can also configure OSPF or static routes according to your network requirements. The PIM, MLT, VRRP, and IGMP settings provided remain unaffected by the choice of routing protocol.

The multicast group range is 224.10.10.0 255.255.255.0.

The STG, MLT, and VLAN number information are displayed in the following figure which shows a sample topology using PIM-SM.

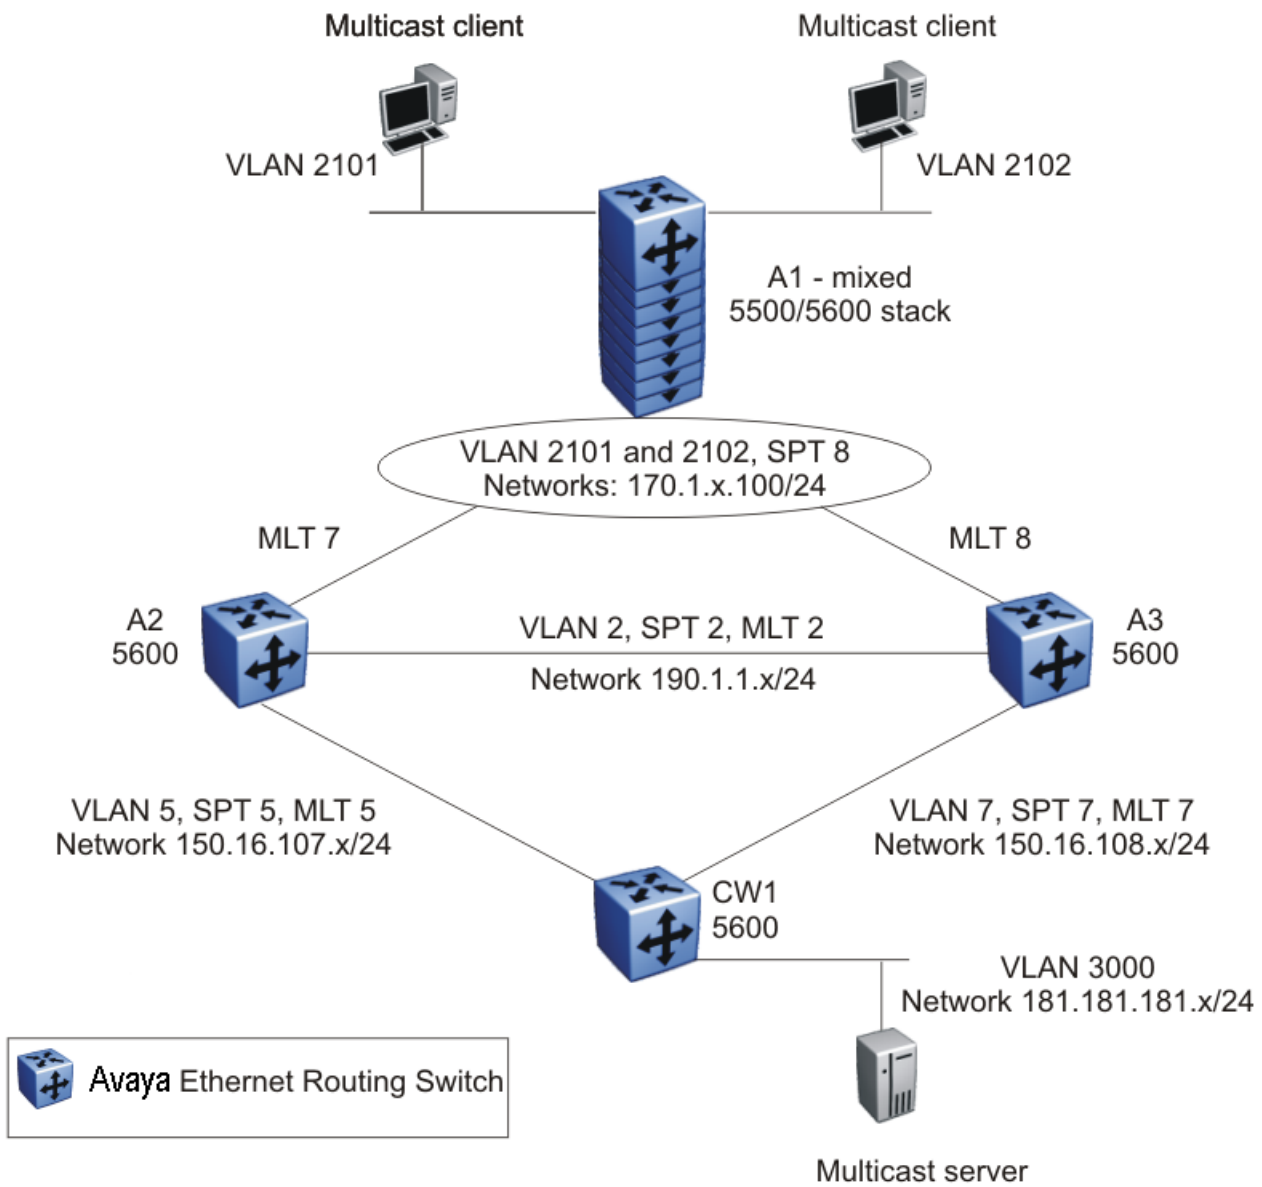


Figure 49: PIM-SM sample topology

A1 description

A1 is an 8-unit mixed stack of 5500 and 5600 Series switches running IGMPv2 snooping. There are a total of two multicast clients on the access layer connected to the A1 stack, each in a different VLAN (2101 and 2102) and in a different network.

For simplicity, the configuration shows only two clients connected to the access layer stack. You can add more ports to each VLAN on the stack to have more users per VLAN.

A2 and A3 description

The distribution layer switches (A2 and A3) are configured as dynamic C-RPs or static RPs (configurations for both options are provided). You can use static RP or dynamic RP (but not both) in accordance with the requirements of your network. If you choose static RP, you must configure the same static RP on every PIM router in your network.

VRRP is enabled on A2 and A3, and all multicast clients have the VRRP virtual IP address as the default gateway for a specific VLAN.

Important:

The VRRP configuration shown is an optional configuration providing a virtual IP for the host gateway. If your network does not need a virtual IP for a gateway, you do not need to configure VRRP. PIM-SM is independent of VRRP.

In this example, A3 is the DR for both PIM client VLANs (2101 and 2102), so all (S,G) entries install on A3. However, you can manage the DR election for the client VLANs by manipulating the IP address of the A2 and A3 VLAN interfaces. To load-share between A2 and A3, you can configure one of the VLAN interfaces on A2 (for example, 2101) with a higher IP address than the corresponding VLAN interface on A3. For the second VLAN, 2102, you can maintain the higher IP address on the A3 interface. In this way, A2 can become the DR for VLAN 2101, and A3 can remain the DR for VLAN 2102. This allows the (S,G) load to be split between the two switches and the system to be used to its maximum limits.

CW1 description

CW1 is configured as the BSR with priority 10 (only applicable to dynamic RP). A higher priority indicates a higher probability of being elected the BSR.

CW1 is directly connected to the multicast server. If desired, you can have a Layer 2 switch between the CW1 and the server with VLAN 3000 spanning through the switch to maintain the connection.

The CW1 connection to the multicast server is configured as a passive interface as it is on the edge and is not required to form a neighbor relationship with any other PIM router. You can configure this interface as an active interface according to the requirements of your network.

Link descriptions

The link connections (port numbers) between devices are listed below; the physical connections are in a one-to-one mapping in sequence as listed for each set of connections.

- A2 – A1:
 - 12,24,36,48,60,72,86,90 – 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2
 - MLT 7, VLAN 2101 to 2128, STG 8
- A3 – A1:
 - 2,14,26,38,50,62,74,80 -- 1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
 - MLT 8, VLAN 2101 to 2128, STG 8
- A2 – A3:
 - 95,96 – 95,96
 - MLT 2, VLAN 2, STG 2
- A2 – CW1:
 - 91,92 – 23,24
 - MLT 5, VLAN 5, STG 5
- A3 – CW1:
 - 91,92 – 21,22
 - MLT 7, VLAN 7, STG 7
- CW1 – Multicast server NIC:
 - 12 – Multicast server NIC
- A1 – Multicast client NICs:
 - VLAN 2101: 1/11 – MC1
 - VLAN 2102: 2/11 – MC2

See the following sections to configure the topology shown. In addition to the listed configurations, you can also configure the optional PIM-SM global and interface parameters, although it is advisable to leave these parameters at their default values.

A1 configuration

The following procedure shows the configuration required for the A1 mixed stack running IGMP snooping.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable tagging on ports:

```
vlan port
1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
tagging enable
```

3. Create the spanning tree instance:

```
spanning-tree stp 8 create
```

4. Configure the VLANs:

```
vlan members remove 1/2,2/14,3/14,4/38,5/12,6/14,7/2,
8/2,1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14

vlan create 2101 type port
vlan members add 2101 1/2,2/14,3/14,4/38,5/12,6/14,7/2,
8/2,1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14,1/11
spanning-tree stp 8 add-vlan 2101
int vlan 2101
ip igmp snooping
ip igmp mrouter 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,
1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14

vlan create 2102 type port
vlan members add 2102 1/2,2/14,3/14,4/38,5/12,6/14,7/2,
8/2,1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14,2/11
spanning-tree stp 8 add-vlan 2102
int vlan 2102
ip igmp snooping
ip igmp mrouter 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,
1/48,2/48,3/2,4/2, 5/14,6/38,7/14,8/14
```

5. Enable spanning tree:

```
spanning-tree 8 enable
```

6. Configure the MLTs:

```
mlt 7 member 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2
mlt 7 enable
```

```
mlt 8 member 1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
mlt 8 enable
```

A2 configuration

The following procedure shows the configuration required for the A2 PIM-SM-enabled distribution layer 5600 Series switch running VRRP and RIP.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable RIP and PIM:

```
ip routing
router rip enable
ip pim enable
```

3. Enable tagging on ports:

```
vlan port 95-96,98,91-92,12,24,36,48,60,72,86,90 tagging enable
```

4. Configure the VLANs:

```
vlan members remove 1 12,24,36,48,60,72,86,90,95,96,91,
92
vlan create 2 type port
vlan members remove 1 95-96
vlan members add 2 95-96
interface vlan 2
ip address 190.1.1.2 255.255.255.0
ip pim en
ip rip en

vlan create 5 type port
vlan members remove 1 91-92
vlan members add 5 91-92
interface vlan 5
ip address 150.16.107.2 255.255.255.0
ip pim en
ip rip en

vlan create 2101 type port
vlan members add 2101 12,24,36,48,60,72,86,90
interface vlan 2101
ip address 170.1.1.1 255.255.255.0
ip pim en
ip rip en
```

5. Configure spanning tree:

```
spanning-tree stp 5 create
spanning-tree stp 5 add-vlan 5
spanning-tree stp 5 enable

spanning-tree stp 2 create
spanning-tree stp 2 add-vlan 2
spanning-tree stp 2 enable

spanning-tree stp 8 create
```

```
spanning-tree stp 8 add-vlan 2101
spanning-tree stp 8 add-vlan 2102
spanning-tree stp 8 enable
```

6. Configure the MLTs:

```
mlt 5 member 91-92
mlt 5 enable
mlt 2 member 95-96
mlt 2 enable
mlt 7 member 12,24,36,48,60,72,86,90
mlt 7 enable
```

7. Configure VRRP:

```
router vrrp enable
interface vlan 2101
ip vrrp add 21 170.1.1.100
ip vrrp 21 enable

interface vlan 2102
ip vrrp add 22 170.1.2.100
ip vrrp 22 enable
```

8. For PIM-SM, configure a static RP:

```
ip pim static-rp enable
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.107.2
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.108.3
```

OR

configure a dynamic C-RP:

```
ip pim rp-candidate group 224.10.10.0 255.255.255.0 rp 150.16.107.2
```

A3 configuration

The following procedure shows the configuration required for the A3 PIM-SM-enabled distribution layer 5600 Series switch running VRRP and RIP.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable RIP and PIM:

```
ip routing
router rip enable
ip pim enable
```

3. Enable tagging on ports:

```
vlan port 95-96,98,91-92,2,14,26,38,50,62,74,80
tagging ena
```

4. Configure the VLANs:

```
vlan members remove 1 2,14,26,38,50,62,74,80
vlan create 2 type port
```

PIM-SM/SSM configuration example using ACLI

```
vlan members remove 1 95-96
vlan members add 2 95-96
interface vlan 2
ip address 190.1.1.3 255.255.255.0
ip pim en
ip rip en

vlan create 7 type port
vlan members remove 1 91-92
vlan members add 7 91-92
interface vlan 7
ip address 150.16.108.3 255.255.255.0
ip pim en
ip rip en

vlan create 2101 type port
vlan members add 2101 2,14,26,38,50,62,74,80
interface vlan 2101
ip address 170.1.1.2 255.255.255.0
ip pim en
ip rip en

vlan create 2102 type port
vlan members add 2102 2,14,26,38,50,62,74,80,49
interface vlan 2102
ip address 170.1.2.2 255.255.255.0
ip pim en
ip rip en
```

5. Configure spanning tree:

```
spanning-tree stp 7 create
spanning-tree stp 7 add-vlan 7
spanning-tree stp 7 enable
spanning-tree stp 2 create
spanning-tree stp 2 add-vlan 2
spanning-tree stp 2 enable
spanning-tree stp 8 create
spanning-tree stp 8 add-vlan 2101
spanning-tree stp 8 add-vlan 2102
spanning-tree stp 8 enable
```

6. Configure the MLTs:

```
mlt 7 member 91-92
mlt 7 enable

mlt 2 member 95-96
mlt 2 enable

mlt 8 member 2,14,26,38,50,62,74,80
mlt 8 enable
```

7. Configure VRRP:

```
router vrrp enable
interface vlan 2101
ip vrrp add 21 170.1.1.100
ip vrrp 21 enab

interface vlan 2102
```

```
ip vrrp add 22 170.1.2.100
ip vrrp 22 enab
```

8. For PIM-SM, configure a static RP:

```
ip pim static-rp enable
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.107.2
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.108.3
```

OR

configure a dynamic C-RP:

```
ip pim rp-candidate group 224.10.10.0 255.255.255.0 rp 150.16.108.3
```

CW1

The following procedure shows the configuration required for the CW1 PIM-SM-enabled 5600 Series switch running RIP. This is the source DR.

The following procedure shows the configuration required for the CW1 PIM-SM-enabled switch running RIP.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable RIP and PIM:

```
ip routing
router rip enable
ip pim enable
```

3. Enable tagging on ports:

```
vlan port 1-2,13-14,21-24,25,29,32 tagging ena
```

4. Configure the VLAN:

```
vlan mem remove 1 23,24,21,22,12

vlan create 5 type port
vlan members add 5 23-24
interface vlan 5
ip address 150.16.107.1 255.255.255.0
ip pim en
ip rip en
```

```
vlan create 7 type port
vlan members add 7 21-22
interface vlan 7
ip address 150.16.108.1 255.255.255.0
ip pim en
ip rip en
```

```
!! THE FOLLOWING VLAN IS A PASSIVE PIM VLAN (As it is connected to the
multicast server, and it does not need to be part of PIM control messages.)
```

```
!! IT CAN BE MADE ACTIVE AS PER YOUR NETWORK
vlan create 3000 type port
```

```
vlan members add 3000 12
interface vlan 3000
ip address 181.181.181.100 255.255.255.0
ip pim interface-type passive
ip pim en
ip rip en
```

5. Configure spanning tree:

```
spanning-tree stp 5 create
spanning-tree stp 5 add-vlan 5
spanning-tree stp 5 enable

spanning-tree stp 7 create
spanning-tree stp 7 add-vlan 7
spanning-tree stp 7 enable
```

6. Configure the MLTs:

```
mlt 5 member 23-24
mlt 5 enable

mlt 7 member 21-22
mlt 7 enable
```

7. For PIM-SM, configure a static RP:

```
ip pim static-rp enable
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.107.2
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.108.3
```

OR

for dynamic RP, configure the C-BSR:

```
interface vlan 5 ip pim bsr-candidate priority 10 exit
```

Example 2

In this second example, A1 is an 8-unit mixed stack of ERS 5500 and ERS 5600 Series switches running IGMPv3 snooping.

A2, A3, and CW1 are all ERS 5600 Series switches with PIM-SSM enabled.

RIP is used as the Layer 3 routing protocol but you can also configure OSPF or static routes according to your network requirements. The PIM, MLT, VRRP, and IGMP settings provided remain unaffected by the choice of routing protocol.

The multicast group range is 224.10.10.0 255.255.255.0.

The STG, MLT, and VLAN number information are displayed in the following figure which shows a sample topology using PIM-SSM.

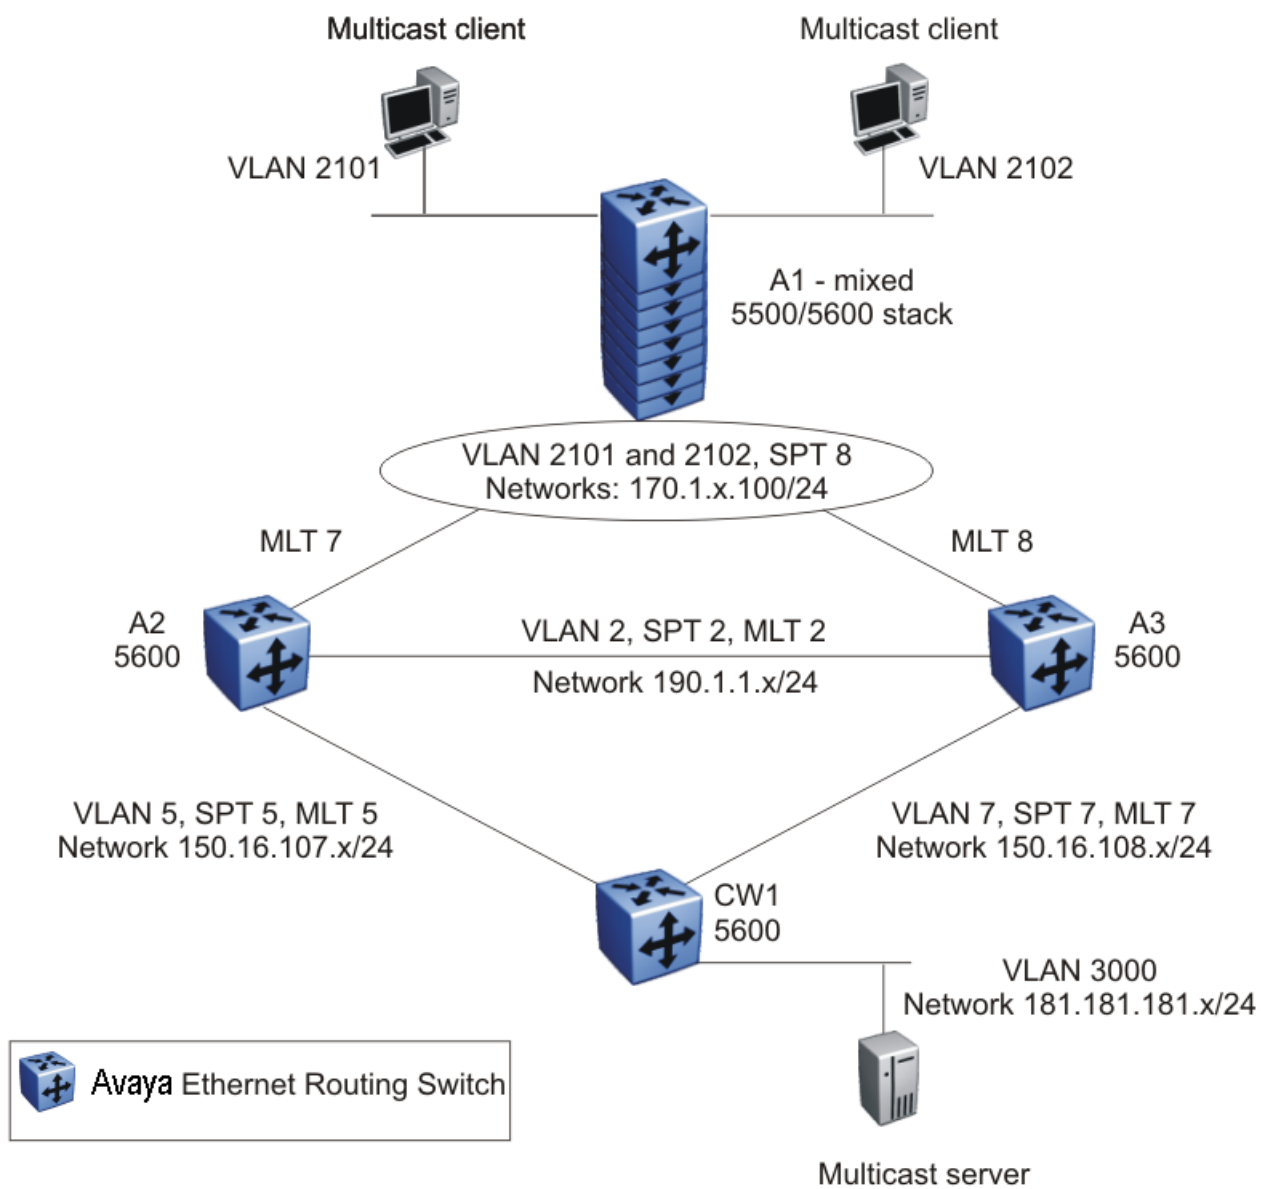


Figure 50: PIM-SSM sample topology

A1 description

A1 is an 8-unit mixed stack of 5500 and 5600 Series switches running IGMPv3 snooping. There are a total of two multicast clients on the access layer connected to the A1 stack, each in a different VLAN (2101 and 2102) and in a different network.

For simplicity, the configuration shows only two clients connected to the access layer stack. You can add more ports to each VLAN on the stack to have more users per VLAN.

A2 and A3 description

The distribution layer switches (A2 and A3) are configured with PIM-SSM.

VRRP is enabled on A2 and A3, and all multicast clients have the VRRP virtual IP address as the default gateway for a specific VLAN.

Important:

The VRRP configuration shown is an optional configuration providing a virtual IP for the host gateway. If your network does not need a virtual IP for a gateway, you do not need to configure VRRP. PIM-SSM is independent of VRRP.

In this example, A3 is the DR for both PIM client VLANs (2101 and 2102), so all (S,G) entries install on A3. However, you can manage the DR election for the client VLANs by manipulating the IP address of the A2 and A3 VLAN interfaces. To load-share between A2 and A3, you can configure one of the VLAN interfaces on A2 (for example, 2101) with a higher IP address than the corresponding VLAN interface on A3. For the second VLAN, 2102, you can maintain the higher IP address on the A3 interface. In this way, A2 can become the DR for VLAN 2101, and A3 can remain the DR for VLAN 2102. This allows the (S,G) load to be split between the two switches and the system to be used to its maximum limits.

CW1 description

CW1 is directly connected to the multicast server. If desired, you can have a Layer 2 switch between the CW1 and the server with VLAN 3000 spanning through the switch to maintain the connection

The CW1 connection to the multicast server is configured as a passive interface as it is on the edge and is not required to form a neighbor relationship with any other PIM router. You can configure this interface as an active interface according to the requirements of your network.

Link descriptions

The link connections (port numbers) between devices are listed below; the physical connections are in a one-to-one mapping in sequence as listed for each set of connections.

- A2 – A1:
 - 12,24,36,48,60,72,86,90 – 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2
 - MLT 7, VLAN 2101 to 2128, STG 8
- A3 – A1:
 - 2,14,26,38,50,62,74,80 -- 1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
 - MLT 8, VLAN 2101 to 2128, STG 8
- A2 – A3:
 - 95,96 – 95,96
 - MLT 2, VLAN 2, STG 2
- A2 – CW1:
 - 91,92 – 23,24
 - MLT 5, VLAN 5, STG 5
- A3 – CW1:
 - 91,92 – 21,22
 - MLT 7, VLAN 7, STG 7
- CW1 – Multicast server NIC:
 - 12 – Multicast server NIC
- A1 – Multicast client NICs:
 - VLAN 2101: 1/11 – MC1
 - VLAN 2102: 2/11 – MC2

A1 configuration

The following procedure shows the configuration required for the A1 mixed stack running IGMP snooping.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable tagging on ports:

```
vlan port
1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
4 tagging enable
```

3. Create the spanning tree instance:

```
spanning-tree stp 8 create
```

4. Configure the VLANs:

```
vlan members remove 11/2,2/14,3/14,4/38,5/12,6/14,7/2,
8/2,1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14

vlan create 2101 type port
vlan members add 2101 1/2,2/14,3/14,4/38,5/12,6/14,7/2,
8/2,1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14,1/11
spanning-tree stp 8 add-vlan 2101
int vlan 2101
ip igmp snooping version 3
ip igmp mrouter 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,
1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14

vlan create 2102 type port
vlan members add 2102 1/2,2/14,3/14,4/38,5/12,6/14,7/2,
8/2,1/48,2/48,3/2,4/2,
5/14,6/38,7/14,8/14,2/11
spanning-tree stp 8 add-vlan 2102
int vlan 2102
ip igmp snooping version 3
ip igmp mrouter 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,
1/48,2/48,3/2,4/2, 5/14,6/38,7/14,8/14
```

5. Enable spanning tree:

```
spanning-tree 8 enable
```

6. Configure the MLTs:

```
mlt 7 member 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2
mlt 7 enable
mlt 8 member 1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
mlt 8 enable
```

A2 configuration

The following procedure shows the configuration required for the A2 PIM-SSM-enabled distribution layer 5600 Series switch running VRRP and RIP.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable RIP and PIM-SSM:

```
ip routing
router rip enable
ip pim enable mode ssm
```

3. Enable tagging on ports:

```
vlan port 95-96,98,91-92,12,24,36,48,60,72,86,90 tagging enable
```

4. Configure the VLANs:

```
vlan members remove 1 12,24,36,48,60,72,86,90,95,96,91,
92
vlan create 2 type port
vlan members remove 1 95-96
vlan members add 2 95-96
interface vlan 2
ip address 190.1.1.2 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

vlan create 5 type port
vlan members remove 1 91-92
vlan members add 5 91-92
interface vlan 5
ip address 150.16.107.2 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

vlan create 2101 type port
vlan members add 2101 12,24,36,48,60,72,86,90
interface vlan 2101
ip address 170.1.1.1 255.255.255.0
ip pim en
ip igmp version 3
ip rip en
```

5. Configure spanning tree:

```
spanning-tree stp 5 create
spanning-tree stp 5 add-vlan 5
spanning-tree stp 5 enable

spanning-tree stp 2 create
spanning-tree stp 2 add-vlan 2
spanning-tree stp 2 enable

spanning-tree stp 8 create
spanning-tree stp 8 add-vlan 2101
spanning-tree stp 8 add-vlan 2102
spanning-tree stp 8 enable
```

6. Configure the MLTs:

```
mlt 5 member 91-92
mlt 5 enable
mlt 2 member 95-96
mlt 2 enable
```

```
mlt 7 member 12,24,36,48,60,72,86,90
mlt 7 enable
```

7. Configure VRRP:

```
router vrrp enable
interface vlan 2101
ip vrrp add 21 170.1.1.100
ip vrrp 21 enable

interface vlan 2102
ip vrrp add 22 170.1.2.100
ip vrrp 22 enable
```

A3 configuration

The following procedure shows the configuration required for the A3 PIM-SSM-enabled distribution layer 5600 Series switch running VRRP and RIP.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable RIP and PIM-SSM:

```
ip routing
router rip enable
ip pim enable mode ssm
```

3. Enable tagging on ports:

```
vlan port 95-96,98,91-92,2,14,26,38,50,62,74,80
tagging ena
```

4. Configure the VLANs:

```
vlan members remove 1 2,14,26,38,50,62,74,80

vlan create 2 type port
vlan members remove 1 95-96
vlan members add 2 95-96
interface vlan 2
ip address 190.1.1.3 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

vlan create 7 type port
vlan members remove 1 91-92
vlan members add 7 91-92
interface vlan 7
ip address 150.16.108.3 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

vlan create 2101 type port
vlan members add 2101 2,14,26,38,50,62,74,80
interface vlan 2101
ip address 170.1.1.2 255.255.255.0
```

```

ip pim en
ip igmp version 3
ip rip en

vlan create 2102 type port
vlan members add 2102 2,14,26,38,50,62,74,80,49
interface vlan 2102
ip address 170.1.2.2 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

```

5. Configure spanning tree:

```

spanning-tree stp 7 create
spanning-tree stp 7 add-vlan 7
spanning-tree stp 7 enable
spanning-tree stp 2 create
spanning-tree stp 2 add-vlan 2
spanning-tree stp 2 enable
spanning-tree stp 8 create
spanning-tree stp 8 add-vlan 2101
spanning-tree stp 8 add-vlan 2102
spanning-tree stp 8 enable

```

6. Configure the MLTs:

```

mlt 7 member 91-92
mlt 7 enable

mlt 2 member 95-96
mlt 2 enable

mlt 8 member 2,14,26,38,50,62,74,80
mlt 8 enable

```

7. Configure VRRP:

```

router vrrp enable
interface vlan 2101
ip vrrp add 21 170.1.1.100
ip vrrp 21 enable

interface vlan 2102
ip vrrp add 22 170.1.2.100
ip vrrp 22 enable

```

CW1

The following procedure shows the configuration required for the CW1 PIM-SSM-enabled 5600 Series switch running RIP. This is the source DR.

The following procedure shows the configuration required for the CW1 PIM-SSM-enabled switch running RIP.

1. Enter Global Configuration mode:

```
configure terminal
```

2. Enable RIP and PIM-SSM:

```
ip routing
router rip enable
ip pim enable mode ssm
```

3. Enable tagging on ports:

```
vlan port 1-2,13-14,21-24,25,29,32 tagging ena
```

4. Configure the VLAN:

```
vlan mem remove 1 23,24,21,22,12

vlan create 5 type port
vlan members add 5 23-24
interface vlan 5
ip address 150.16.107.1 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

vlan create 7 type port
vlan members add 7 21-22
interface vlan 7
ip address 150.16.108.1 255.255.255.0
ip pim en
ip igmp version 3
ip rip en

!! THE FOLLOWING VLAN IS A PASSIVE PIM VLAN (As it is connected to the
multicast server, and it does not need to be part of PIM control messages.)

!! IT CAN BE MADE ACTIVE AS PER YOUR NETWORK
vlan create 3000 type port
vlan members add 3000 12
interface vlan 3000
ip address 181.181.181.100 255.255.255.0
ip pim interface-type passive
ip pim en
ip igmp version 3
ip rip en
```

5. Configure spanning tree:

```
spanning-tree stp 5 create
spanning-tree stp 5 add-vlan 5
spanning-tree stp 5 enable

spanning-tree stp 7 create
spanning-tree stp 7 add-vlan 7
spanning-tree stp 7 enable
```

6. Configure the MLTs:

```
mlt 5 member 23-24
mlt 5 enable

mlt 7 member 21-22
mlt 7 enable
```

Chapter 39: PIM-SM/SSM configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure PIM-SM/SSM.

Unlike dense-mode protocols, such as Distance Vector Multicast Routing Protocol (DVMRP), that initially flood multicast traffic to all routers over an entire internetwork, PIM sends multicast traffic only to routers that belong to a specific multicast group and that choose to receive the traffic. PIM reduces overhead costs for processing unwanted multicast packets.

Prerequisites

Before you can configure PIM, you must prepare the switch as follows:

1. Install the Advanced Routing software license.

! Important:

If your Ethernet Routing Switch is running an Advanced License for a release prior to Release 6.0, to enable PIM you must regenerate your license file from the Avaya web site and install the new license file on the switch.

2. Configure IP addresses on the VLAN interfaces on which you want to configure PIM.
3. Enable a unicast protocol, either RIP or OSPF, globally and on the interfaces on which you want to configure PIM.

! Important:

PIM requires a unicast protocol to multicast traffic within the network when performing the Reverse Path Forwarding (RPF) check. PIM also uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree. The unicast routing table must contain a route to every multicast source in the network, as well as routes to PIM entities such as the rendezvous points (RP) and bootstrap router (BSR).

Configuring PIM-SM/SSM

Use the following procedure to configure PIM-SM.

1. Enable PIM globally.
(If desired, modify the default global PIM properties.)
2. Enable PIM on individual VLAN interfaces.
(If desired, modify the default VLAN PIM properties.)
3. For PIM-SM, configure candidate RPs for the multicast groups in the network. (It is best to have multiple candidate-RPs in the network; however, with the Ethernet Routing Switch 5000, you can only configure one candidate-RP per switch for any number of groups.)

OR

Configure one (or several) static RPs for the multicast groups in the network. (To enable static RP in the PIM-SM domain, you must configure the same static RPs on every system that takes part in PIM forwarding.)

4. For PIM-SM, configure one or several candidate BSRs to propagate RP information to all switches in the network. You can configure every PIM-enabled VLAN as a C-BSR. (If Static RP is enabled, this step is not required.)

! **Important:**

Ensure that all routers in the path from the receivers to the RP and to the multicast source are PIM-enabled. Also ensure that all PIM routers have unicast routes to reach the source and RP through directly-connected PIM neighbors.

Use the following procedure to configure PIM-SSM.

1. Enable PIM globally and change PIM mode to SSM.
(If desired, modify the default global PIM properties.)
2. Enable PIM on individual VLAN interfaces.
(If desired, modify the default VLAN PIM properties.)
3. If you use PIM-SSM with the IGMPv3 protocol, then configure this option on each VLAN.

All additional configurations listed below are optional and can be configured according to the requirements of your network.

Configuring global PIM-SM/SSM status and properties

Use the following procedure to configure PIM-SM/SSM status and properties globally.

Prerequisites

- You must enable PIM-SM globally.

Procedure steps

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **PIM**.
3. Select the **Globals** tab.
4. In the Globals tab, in the **Mode** box, select **sm** for PIM-SM or **ssm** for PIM-SSM.
5. Select the **Enable** check box to enable PIM-SM/SSM.
6. Configure the other parameters as required.
7. In the toolbar, click **Apply**.

Variable definitions

Field	Description
Mode	Displays the PIM mode on the switch: sparse mode or source specific multicast mode.
Enable	Enables or disables PIM globally.
JoinPruneInterval	Specifies how long (in seconds) the PIM router waits between sending join/prune messages to its upstream neighbors. The range is 1–18724, and the default is 60 seconds.
RegisterSuppTimer	Specifies how long (in seconds) the DR suppresses sending register messages to the RP after the DR receives a register-stop message from the RP. The range is 5–65535, and the default is 60 seconds.

Field	Description
UniRouteChgTimeOut	<p>Specifies how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates to be used by PIM. The range is 2–65535, and the default is 5 seconds.</p> <p>! Important: Lowering this value increases how often the switch polls the RTM. This can affect the performance of the switch, especially when a high volume of traffic is flowing through the switch.</p>
DiscardDataTimeOut	<p>After the router forwards the first source packet to the RP, this value specifies how long (in seconds) the router discards subsequent source data while waiting for a join from the RP. An IPMC discard record is created and deleted after the timer expires or after a join is received. The range is 5–65535, and the default is 60 seconds.</p>
CRPADVTimeOut	<p>Specifies how often (in seconds) routers that are configured as candidate RPs send candidate rendezvous point (C-RP) advertisement messages. After this timer expires, the C-RP sends an advertisement message to the elected BSR. The range is 5–26214, and the default is 60 seconds.</p>
BootStrapPeriod	<p>Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. The range is 5–32757, and the default is 60 seconds.</p>
StaticRP	<p>Enables or disables the static RP feature. Static RP permits communication with routers from other vendors that do not use the BSR mechanism. By default, static RP is disabled.</p>
FwdCacheTimeOut	<p>Specifies the PIM forward cache expiry value in seconds. This value is used in aging PIM mroutes in seconds. The range is 10–86400, and the default is 210.</p>

Configuring PIM-SM/SSM status and properties for a VLAN

Use the following procedure to enable PIM on a VLAN and configure related properties.

By default, PIM-SM is disabled on VLANs.

Prerequisites

- You must enable PIM-SM globally.
- Before you change the state (active or passive) of a PIM interface using the InterfaceType field, first disable PIM to prevent instability in the PIM operations, especially when neighbors are present or when streams are received.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs**.
3. In the **Basic** tab, select the VLAN ID that you want to configure with PIM.
4. In the toolbar, click **IP**.
5. In the work area, click the **PIM** tab.
6. Select the **Enable** check box.
7. Configure the parameters as required.
8. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **PIM** tab.

Field	Description
Enable	Enables or disables PIM-SM on the VLAN.
Mode	Displays the PIM mode on the switch: sparse mode or source specific multicast mode.
HelloInterval	Specifies the interval (in seconds) that the PIM router waits between sending out hello message to neighboring routers. The default is 30 seconds.
JoinPruneInterval	Specifies the interval (in seconds) the PIM router waits between sending out join/prune message to its upstream neighbors. The default is 60 seconds.
CBSRPreference	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR.

Field	Description
InterfaceType	Specifies the state (active or passive) of PIM on a VLAN interface. An active interface transmits and receives PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. Passive interfaces are useful when you have a high number of PIM interfaces and these interfaces are connected to end users, not to other switches. By default, PIM-SM interfaces are active.

Configuring PIM-SM/SSM VLAN properties from the IP Routing menu

After you have enabled PIM on a VLAN, use the following procedure to view and edit PIM VLAN parameters from the PIM Interfaces tab accessible under the IP Routing menu. This procedure does not provide more configuration options than are available under the VLAN menu, but it does allow you to view some additional PIM parameters (such as DR) and also to view the configuration for multiple VLANs at once.

Prerequisites

- You must enable PIM-SM globally.
- You must enable PIM-SM on a VLAN.
- Before you change the state (active or passive) of a PIM interface using the Interface Type field, first disable PIM to prevent instability in the PIM operations, especially when neighbors are present or when streams are received.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Interfaces** tab.
4. In the table, double-click the cell under the column heading for the parameter you want to change.
5. Select a parameter or value from the drop-down list.
6. Repeat the previous two steps until you have amended all of the parameters you want to change.
7. In the toolbar, click **Apply**.

Variable definitions

The following table describes the fields of the **Interfaces** tab.

Field	Description
IfIndex	Specifies the VLANs configured for PIM-SM.
Address	Specifies the IP address of the PIM-SM VLAN.
NetMask	Specifies the network mask for the PIM-SM VLAN.
Enable	Specifies the status of PIM-SM on the VLAN: enabled (true) or disabled (false).
Mode	Specifies the PIM mode: sparse mode or source specific multicast mode.
DesignatedRouter	Specifies the router with the highest IP address on a LAN designated to perform the DR tasks.
HelloInterval	Specifies the interval (in seconds) the switch waits between sending hello message to neighboring switches. The default is 30 seconds.
JoinPruneInterval	Specifies the interval (in seconds) the switch waits between sending join/prune message to its upstream neighbors. The default is 60 seconds.
CBSRPreference	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR-priority and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR.
InterfaceType	Specifies the type of interface: active or passive.
OperState	Indicates the operating status of PIM-SM on this interface: up or down.

Specifying the router as a candidate BSR on a VLAN interface using EDM

Because PIM-SM cannot run without a bootstrap router (BSR), you must specify at least one C-BSR in the domain. Any additional C-BSRs provide backup protection in case the primary BSR fails.

The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs have equal priority, the candidate with the higher IP address becomes the BSR. If you

add a new C-BSR with the highest priority to the domain, it automatically becomes the new BSR.

With the Ethernet Routing Switch 5000 Series, you can configure every PIM-enabled interface as a C-BSR.

Setting the C-BSR priority using EDM from the VLAN menu

Use the following procedure to set the C-BSR priority on a VLAN from the VLAN menu.

Procedure steps

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, double-click **VLANs**.
3. In the **Basic** tab, select the VLAN ID that you want to configure with PIM.
4. In the toolbar, click **IP**.

The **IP, VLAN** tab appears.

5. In the work area, click the **PIM** tab.
6. In the **CBSRPreference** field, type the value of the C-BSR priority.
7. In the toolbar, click **Apply**.

Setting the C-BSR priority using EDM from the IP Routing menu

Use the following procedure to set the C-BSR priority on a VLAN from the IP Routing menu.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Interfaces** tab.
4. In the table, double-click the cell under the CBSRPreference column heading for the parameter you want to change.
5. Type the value of the C-BSR priority for the associated interface.

The Candidate BSR with the highest BSR-priority and address is the preferred BSR. The default is -1 , which indicates that the current interface is not a Candidate BSR; the range is 0 to 255.

6. In the toolbar, click **Apply**.

Displaying the current BSR using EDM

Use the following procedure to display the current BSR information to review the configuration.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Current BSR** tab to view the current BSR information.

Variable definitions

The following table describes the fields of the **Current BSR** tab.

Field	Description
Address	Specifies the IP address of the current BSR for the local PIM domain.
FragmentTag	Specifies a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same Fragment Tag.
HashMask	Specifies the mask used in the hash function to map a group to one of the C-RPs from the RP set. With the hash mask, a small number of consecutive groups can always hash to the same RP.
Priority	Specifies the priority of the current BSR. The candidate BSR (C-BSR) with the highest BSR priority, and address (referred to as the preferred BSR) is elected as the BSR for the domain.
BootStrapTimer	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages.

Specifying a local IP interface as a candidate RP using EDM

Because PIM-SM cannot run without an RP, you must specify at least one C-RP in the domain.

With the Ethernet Routing Switch 5000 Series, you can configure only one local interface as a C-RP for any number of groups.

Using the GroupMask value, you can configure a C-RP for several groups in one configuration. For example, with a C-RP configuration with a GroupAddress value of 224.0.0.0 and a GroupMask of 240.0.0.0, you can configure the C-RP for a multicast range from 224.0.0.0 to 239.255.255.255.

Use the following procedure to configure a local PIM-SM interface as a candidate RP (C-RP).

Variable definitions

The following table describes the fields of the Insert Candidate RP dialog box.

Field	Description
GroupAddress	Specifies the IP address of the multicast group. Together with the group mask, the group address identifies the prefix that the local router uses to advertise itself as a C-RP.
GroupMask	Specifies the address mask of the multicast group. Together with the group address, the group mask identifies the prefix that the local router uses to advertise itself as a C-RP.
RPAddress	Specifies the IP address of the C-RP. This address must be one of the local PIM-SM enabled interfaces.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Candidate RP** tab.
4. In the toolbar, click **Insert**.

The Insert Candidate RP dialog box appears.

5. In the **GroupAddress** box, enter the multicast group address.
6. In the **GroupMask** box, enter the multicast group mask.
7. In the **RPAddress** box, enter the address of the C-RP.
8. Click **Insert**.

Displaying the active RP using EDM

Use the following procedure to display the active RP.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Active RP** tab.

Variable definitions

The following table describes the Active RP dialog box fields.

Field	Description
GroupAddress	Specifies the IP address of the multicast group.
GroupMask	Specifies the address mask of the multicast group.
ActiveRP	Specifies the IP address of the active RP.
Priority	Specifies the priority of the active RP.

Configuring a static RP using EDM

Use the following procedure to configure a static RP entry. After you configure static RP, the switch ignores the BSR mechanism and uses the statically-configured RPs only.

Prerequisites

- You must enable static RP.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Static RP** tab.
4. In the toolbar, click **Insert**.

The Insert Static RP dialog box appears.

5. In the **GroupAddress** box, type the multicast group address.
6. In the **GroupMask** box, type the multicast group mask.
7. In the **RPAddress** box, enter the address of the static RP.
8. Click **Insert**.

Variable definitions

The following table describes the fields of the **Static RP** tab.

Field	Description
GroupAddress	Specifies the IP address of the multicast group. Together with the group mask, the IP address identifies the range of the multicast addresses that the RP handles.
GroupMask	Specifies the address mask of the multicast group. Together with the group address, the address mask identifies the range of the multicast addresses that the RP handles.
RPAddress	Specifies the IP address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid when the switch has a unicast route to the network for the static RP and is invalid otherwise.

Enabling static RP using EDM

Enable static RP to avoid the process of dynamically learning C-RPs through the BSR mechanism. With this feature, static RP-enabled Ethernet Routing Switch 5000 Series

switches can communicate with switches from other vendors that do not use the BSR mechanism.

Use the following procedure to enable static RP.

! Important:

When you enable static RP, all dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the Globals tab, select the **Enable** check box to enable PIM-SM globally.
4. Select the **StaticRP** check box.
5. In the toolbar, click **Apply**.

Specifying a virtual neighbor on an interface using EDM

Use the following procedure to configure a virtual neighbor when the next hop for a static route cannot run PIM-SM, such as a Virtual Redundancy Router Protocol address on an adjacent device.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Virtual Neighbors** tab.
4. In the toolbar, click **Insert**.
The Insert Virtual Neighbors dialog box appears.
5. In the **NeighborIfIndex** field, click **VLAN**.
6. Select the desired VLAN, and then click **OK**.
7. In the **NeighborAddress** field, enter the IP address of the virtual neighbor.
8. Click **Insert**.

Variable definitions

The following table describes the fields of the **Neighbors** tab.

Field	Description
NeighborIfIndex	Specifies the VLAN ID of the interface used to reach this PIM virtual neighbor.
NeighborAddress	Specifies the IP address of the PIM virtual neighbor.

Displaying PIM-SM/SSM neighbor parameters using EDM

Use the following procedure to view PIM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **Neighbors** tab to view PIM-SM neighbor parameters.

Variable definitions

The following table describes the fields of the **Neighbors** tab.

Field	Description
Address	Specifies the IP address of the PIM neighbor.
IfIndex	Specifies the VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Specifies the elapsed time since this PIM neighbor last became a neighbor of the local router.
ExpiryTime	Specifies the time remaining before this PIM neighbor times out.

Displaying the PIM-SM RP set using EDM

Use the following procedure to display the RP set for troubleshooting purposes. The BSR constructs the RP set from C-RP advertisements and then distributes it to all PIM routers in the PIM domain for the BSR.

Procedure steps

1. From the navigation tree, double-click **IP Routing**.
2. In the IP Routing tree, double-click **PIM**.
3. In the work area, click the **RP Set** tab to view RP Set information.

Variable definitions

The following table describes the fields of the **RP Set** tab.

Field	Description
GroupAddress	Specifies the IP address of the multicast group. Together with the group mask, the group address identifies the prefix that a router uses to advertise itself as a C-RP.
GroupMask	Specifies the address mask of the multicast group. Together with the group address, the group mask identifies the prefix that a router uses to advertise itself as a C-RP.
Address	Specifies the IP address of the C-RP.
HoldTime(sec)	Indicates the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
ExpiryTime	Specifies the time remaining before this C-RP times out.

Chapter 40: Basic IPv6 routing configuration using ACLI

This chapter describes how to use the ACLI to configure basic IPv6 routing.

Basic IPv6 routing configuration procedures

To configure basic IPv6 routing, perform the following steps:

1. Enable global IPv6 routing on the switch.
2. Configure a minimum of one link-local, site-local, or global IPv6 address for each IPv6 VLAN.
3. Configure neighbor discovery prefixes, if desired.
4. Configure static routes, if desired.
5. Configure router advertisement and ICMP properties, as well as the neighbor cache, if desired.

Navigation

- [Configuring global IPv6 routing status](#) on page 528
- [Displaying global IPv6 configuration](#) on page 528
- [Configuring an IPv6 address for a VLAN](#) on page 529
- [Removing the IPv6 address configuration from a VLAN](#) on page 530
- [Displaying IPv6 address configuration for a VLAN](#) on page 531
- [Configuring neighbor discovery prefixes](#) on page 531
- [Displaying neighbor discovery prefix configuration](#) on page 533
- [Configuring router advertisement](#) on page 533
- [Configuring IPv6 ICMP](#) on page 535
- [Configuring IPv6 static routes](#) on page 536
- [Displaying IPv6 static routes](#) on page 537

- [Adding static entries to the neighbor cache](#) on page 538
- [Displaying the neighbor cache](#) on page 539

Configuring global IPv6 routing status

Use this procedure to enable and disable global IPv6 routing at the switch level. By default, IPv6 routing is disabled.

Procedure steps

1. To enable the global IPv6 administrative status, enter the following in Global Configuration mode:

```
[no] ipv6 enable
```
2. To enable IPv6 forwarding, enter the following in Global Configuration mode:

```
[no] ipv6 forwarding
```
3. To configure the IPv6 hop-limit, enter the following in Global Configuration mode:

```
ipv6 hop-limit <hop-limit>
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables the specified parameter.
hop-limit <hop-limit>	Specifies the maximum number of hops before packets drop. Valid range is 0-255.

Displaying global IPv6 configuration

Use the following procedure to display the global IPv6 configuration.

Procedure steps

1. To show the global IPv6 configuration, enter:

```
show ipv6 global
```

2. To show the global IPv6 forwarding status, enter:

```
show ipv6 forwarding
```

Configuring an IPv6 address for a VLAN

Configure an IPv6 address on a VLAN to allow IPv6 routing on the interface.

Procedure steps

1. To configure a link-local identifier, enter the following from the VLAN Interface Configuration mode:

```
[default] ipv6 interface link-local <link-local>
```

2. To configure a site-local or global IPv6 address, enter the following from the VLAN Interface Configuration mode:

```
[no] ipv6 interface address <ipv6 address>
```

3. To configure additional parameters for the IPv6 interface, enter the following from the VLAN Interface Configuration mode:

```
[default] ipv6 interface  
[mtu <bytes>]  
[name <name>]  
[reachacble-time <ms>]  
[retransmit-time <ms>]
```

4. To enable the IPv6 interface, enter the following from the VLAN Interface Configuration mode:

```
[no] [default] ipv6 interface enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified configuration or parameter.
default	Configures the specified parameter to the default value.
address <ipv6 address>	Configures the IPv6 address and prefix length. The default value is none.
link-local <link-local>	Configures the link local identifier. The default value is none.
mtu <bytes>	Configures the maximum transmission unit for the interface. The default value is 1500.
name <name>	Configures a description for the interface. This variable does not support the default parameter.
reachable-time <ms>	Configures the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. Range is 0-3600000. The default value is 30000.
retransmit-time <ms>	Configures the time, in milliseconds, between retransmissions of Neighbor Solicitation messages to a neighbor when resolving the address or when probing the reachability of a neighbor. Range is 0-3600000. The default value is 1000.

Removing the IPv6 address configuration from a VLAN

Use the following procedure to disable the IPv6 interface status and delete the IPv6 address from a VLAN.

Procedure steps

To disable the IPv6 interface status and delete the IPv6 address, enter the following from the VLAN Interface Configuration mode:

```
{no | default} ipv6 interface all
```

Variable definitions

The following table describes the command variables.

Variable	Value
no	Disables the IPv6 interface status and deletes the IPv6 address.
default	Disables the IPv6 interface status and deletes the IPv6 address.

Displaying IPv6 address configuration for a VLAN

Use the following procedure to display the IPv6 address configuration on a VLAN.

Procedure steps

To display the IPv6 interface configuration, enter:

```
show ipv6 interface
```

Configuring neighbor discovery prefixes

Use this procedure to specify the neighbor discovery prefixes to advertise in the router advertisement messages on a VLAN. Configuring prefixes allows for host autoconfiguration of site-local and global IPv6 addresses.

Procedure steps

1. To configure neighbor discovery prefixes, enter the following from the VLAN Interface Configuration mode:

```
ipv6 nd prefix-interface <ipv6address-prefix> [eui <1-3>]
[no-autoconfig <false|true>] [no-advertise] [no-onlink
<false|true>]
```

2. To configure neighbor discovery prefix parameters, enter the following from the VLAN Interface Configuration mode:

```
[no] [default] ipv6 nd prefix <ipv6address/prefix-length>
[infinite] [no-advertise] [preferred-life <0-3600000>]
[valid-life <0-3600000>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes the specified configuration.
[default]	Sets the specified parameter to the default value.
eui <1-3>	Specifies the EUI parameter setting: <ol style="list-style-type: none"> 1. Extended Unique Identifier (EUI) is not used 2. EUI with U/L (Universal/Local bit) complement is enabled 3. EUI is used without U/L The default value is 1 - EUI is not used.
no-advertise	Specifies whether the prefix is advertised. If configured, this parameter prevents prefix advertisement. The default value is false.
no-autoconfig <false true>	If true, the prefix is used for autonomous address configuration. The default value is true.
no-onlink <false true>	If true, onlink determination uses the prefix. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default value is true.
infinite	If configured, the prefix does not expire. The default value is false.
preferred-life <0-3600000>	The number of seconds that the prefix can accept and use new connections. The default value is 604800.
valid-life <0-3600000>	The number of seconds that the prefix advertised in the neighbor advertisement is valid. During the valid lifetime, existing connections can be used. New connections cannot be opened. The default value is 2592000.

Displaying neighbor discovery prefix configuration

Use the following procedure to display the neighbor discovery prefix configuration.

Procedure steps

1. To display the discovery prefix configuration, enter:

```
show ipv6 nd interface [vlan [<vid>]] [details]
```

2. To display the discovery prefixes, enter:

```
show ipv6 nd-prefix interface [vlan [<vid>]] [details]
```

Variable definitions

The following table describes the command variables.

Variable	Value
vlan [<vid>]	Specifies the VLAN for which to display the configuration.
details	Specifies detailed command output.

Configuring router advertisement

Use router advertisement to discover potential default routers in a network and to discover link information.

Procedure steps

To configure router advertisement on a VLAN, enter the following from the VLAN Interface configuration mode:

```
[default] [no] ipv6 nd [dad-ns <0-600>]
[hop-limit <1-255>] [managed-config-flag]
[other-config-flag]
[ra-lifetime <0|4-9000>]
```

```
[rtr-advert-max-interval <4-1800>] [rtr-advert-min-interval
<3-1350>] [send-ra]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Configures the specified parameter to the default value.
[no]	Deletes or disables the specified parameter.
dad-ns	The number of neighbor solicitation messages from duplicate address detection. The acceptable range is 0-600. A value of 0 disables duplicate address detection on the specified interface. A value of 1 configures a single transmission without follow-up transmissions. Use the default operator to configure this value to the default setting. The default value is 1.
hop-limit	Enter the maximum number of hops before packets drop. Use the default operator to configure this value to the default setting. The default value is 30.
managed-config-flag	Configure to true to enable M-bit (managed address configuration) on the router. Use the no operator to remove this option. Use the default operator to configure this value to the default setting. The default value is false.
other-config-flag	Configure to true to enable the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. Use the no operator to remove this option. Use the default operator to configure this value to the default setting. The default value is false.
ra-lifetime	Enter the router lifetime included in router advertisement. Other devices use this information to determine if the router can be reached. The range is 0 or 4–9000. Use the default operator to configure this value to the default setting. The default value is 1800.

Variable	Value
rtr-advert-max-interval	Configures the maximum time allowed between sending unsolicited multicast router advertisements. The default value is 600.
rtr-advert-min-interval	Configures the minimum time allowed, in seconds (3–1350), between sending unsolicited multicast router advertisements from the interface. Use the default operator to configure this value to the default setting. The default value is 200.
send-ra	Enables or disables periodic router advertisement messages. Use the no operator to remove this option. Use the default operator to configure this value to the default setting. The default value is true.

Configuring IPv6 ICMP

Configure Internet Control Message Protocol (ICMP) parameters to modify the transport error and information messages within IPv6 packets.

Procedure steps

1. To configure the ICMP rate, enter the following from Global Configuration mode:


```
[no] [default] ipv6 icmp error-interval <0-2147483647>
```
2. To set the status for redirect messages, enter the following from Global Configuration mode:


```
[no] [default] ipv6 icmp redirect-msg
```
3. To configure the status for unreachable messages, enter using the following from Global Configuration mode:


```
[no] [default] ipv6 icmp unreach-msg
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the configured parameter.
[default]	Sets the specified parameter to the default value.
error-interval <0-2147483647>	Configures the error interval in milliseconds. The interval is the time between transmission of error messages. To configure this option to the default value, use the default operator with the command. The default value is 1000.
redirect-msg	Configures the administrative status for ICMP redirect messages. Use the no operator to remove this configuration. To configure this option to the default value, use the default operator with the command. The default value is disable.
unreach-msg	Configures the administrative status for ICMP unreachable messages. Use the no operator to remove this configuration. To configure this option to the default value, use the default operator with the command. The default value is disable.

Configuring IPv6 static routes

Create a new static route or modify existing static route parameters.

Procedure steps

1. To create the static route, enter the following from the Global Configuration mode:

```
ipv6 route <ipv6address/prefix> next-hop <ipv6address/
prefix>
```

2. To assign a route cost, enter the following from the Global Configuration mode:

```
ipv6 route <ipv6address/prefix> cost <1-65535>
```

3. To configure a route preference, enter the following from the Global Configuration mode:

```
ipv6 route <ipv6address/prefix> preference <1-255>
```

4. To specify an interface used to reach the next-hop, enter the following from the Global Configuration mode:

```
ipv6 route <ipv6address/prefix> [tunnel <tunnel-id>] [vlan
<vlan id>]
```

5. To enable the route, enter the following from the Global Configuration mode:


```
ipv6 route <ipv6address/prefix> enable
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the configured route, or the configured parameter. This option is applicable with the ipv6address/prefix, next-hop, vlan, tunnel, or enable parameters.
[default]	Sets the specified parameter to the default value. This option is applicable with the ipv6address/prefix, cost, enable, next-hop, preference, tunnel, or vlan parameters.
<ipv6address/prefix>	Specifies the IPv6 address and prefix for the route destination as a string of 0–49 characters.
next-hop <ipv6address/prefix>	Specifies the IPv6 address of the next-hop router—the next router at which packets must arrive on this route. The string length is 0–49 characters.
tunnel <tunnel-id>	Specifies the tunnel ID in the range of 1-5000.
vlan <vlan id>	Specifies the VLAN ID in the range of 1–4094.
cost <1-65535>	Specifies the metric of the route in the range of 1–65535.
preference <1-255>	Specifies the route preference in the range of 1–255. The default value is 0.

Displaying IPv6 static routes

Use this procedure to display IPv6 static routes.

Procedure steps

1. To display the static route configuration, enter the following in the PrivExec mode:

```
show ipv6 route static
```

2. To display the route configuration for a particular destination, next-hop, tunnel, or VLAN, enter the following in the PrivExec mode:

```
show ipv6 route
 [dest <ipv6-address/prefix>]
 [next-hop <ipv6-address/prefix>]
```

```
[tunnel <tunnel-id>]
[vlan <vid>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
<ipv6address/prefix>	Displays entries for the specified route destination.
next-hop <ipv6address/prefix>	Displays entries for the specified next-hop address.
tunnel <tunnel-id>	Displays entries for the specified tunnel ID.
vlan <vlan id>	Displays entries for the specified VLAN ID.

Adding static entries to the neighbor cache

The neighbor cache contains information about IPv6 neighbors to which the IPv6 device sends traffic. You can manually add neighbors to the cache.

Procedure steps

To add an entry to the neighbor cache, enter the following from the Global Configuration mode:

```
ipv6 neighbor <ipv6 address> port <port> mac <mac address>
[vlan <vlan id>]
```

Variable definitions

The following table describes the command variables.

<ipv6 address>	Specifies the IPv6 address in hexadecimal colon format {string length 0..128}. The default value is none.
<mac address>	Specifies the MAC address in the following format: {0x00:0x00:0x00:0x00:0x00:0x00}
<port>	Specifies the port on which to add a neighbor.

<vlan id>	Specifies the ID of the VLAN on which to add a neighbor.
-----------	--

Displaying the neighbor cache

Use the following procedure to view entries in the neighbor cache.

Procedure steps

To display entries in the neighbor cache, enter the following from the PrivExec mode:

```
show ipv6 neighbor [<ipv6addr>] [type {other|dynamic|static|
local}] [interface <interface-type> <interface-id>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[<ipv6addr>]	Specifies the neighbor IPv6 address.
[type {other dynamic static local}]	Specifies the type of mapping as one of the following: <ul style="list-style-type: none"> • dynamic: dynamically learned neighbor • local: local neighbor address • other: other neighbor entry • static: manually configured neighbor

Chapter 41: Basic IPv6 routing configuration using Enterprise Device Manager

This chapter describes how to configure basic IPv6 routing using EDM.

Basic IPv6 routing configuration procedures

To configure basic IPv6 routing, perform the following steps:

1. Enable global IPv6 routing on the switch.
2. Configure a minimum of one link-local, site-local, or global IPv6 address for each IPv6 VLAN.
3. Configure neighbor discovery prefixes, if desired.
4. Configure static routes, if desired.
5. Configure router advertisement and ICMP properties, as well as the neighbor cache, if desired.

Configuring IPv6 routing and ICMP

Enable IPv6 routing to route IPv6 traffic on the switch. IPv6 packets transport Internet Control Message Protocol (ICMP) error and information messages. Configure the rate, in milliseconds, at which ICMP sends messages to conserve system resources.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. To enable IPv6 routing, in the **Forwarding** box, select **forwarding**.

4. Configure the routing and ICMP parameters as required.
5. Click **Apply**.

Variable definitions

The following table describes the fields of the **IPv6 Globals** tab.

Variable	Value
AdminEnabled	Configures IPv6 as administratively enabled or disabled.
OperEnabled	Indicates whether IPv6 is operationally enabled or disabled.
Forwarding	Configures whether this entity is an IPv6 router with respect to the forwarding of datagrams received by, but not addressed to, this entity. Select forwarding to act as a router. Select notForwarding to not act as a router. The default is notForwarding.
DefaultHopLimit	Configures the hop limit. The default is 30.
IcmpNetUnreach	If selected, enables the ICMP network unreachable feature. The default is disabled.
IcmpRedirectMsg	If selected, enables the ICMP redirect message feature. The default is disabled.
IcmpErrorInterval	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000 milliseconds. An entry of 0 seconds results in no sent ICMPv6 error messages.
IcmpErrorQuota	Specifies the number of ICMP error messages that can be sent during the ICMP error interval. A value of zero specifies not to send any. The default value is 50.
MulticastAdmin Status	Specifies the multicast administrative status. The default is false.

Configuring a link-local address for a VLAN

Use this procedure to configure an IPv6 link-local address for a VLAN.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.

! **Important:**

Enterprise Device Manager provides multiple paths to configure IPv6 interfaces and addresses. In addition to selecting Configuration, IPv6, IPv6, you can select Configuration, VLAN, VLANs, (select a VLAN), IPv6.

3. Click the **Interfaces** tab.
4. Click **Insert**.
5. In the **IfIndex** box, click **VLAN**, and select a VLAN.
6. You must select the **AdminStatus** check box before the interface takes effect.
7. Edit the remaining fields.
8. Click **Insert**.
9. Click **Apply**.

Variable definitions

The following table describes the fields of the **Interfaces** tab.

Variable	Value
IfIndex	A unique value to identify a logical interface (VLAN). For a VLAN it is the ifindex of the VLAN.
Identifier	The IPv6 address interface identifier. This is a binary string of up to 8 octets in network byte order.
IdentifierLength	The length of the interface identifier in bits.
Descr	A text string containing information about the interface. The network management system also configures this string.
VlanId	A value that uniquely identifies the Virtual LAN associated with the entry. This value corresponds to the lower 12 bits in the IEEE 802.1Q VLAN tag.
Type	The type of interface.

Variable	Value
ReasmMaxSize(MTU)	The MTU for this IPv6 interface. This value must be same for all the IP addresses defined on this interface. The default value is 1500.
PhysAddress	The media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	The indication of whether IPv6 is enabled (true) or disabled (false) on this interface. This object does not affect the state of the interface itself, only the connection to an IPv6 stack. The default is false.
OperStatus	The current operational status of the interface.
ReachableTime	The time (in milliseconds) a neighbor is considered reachable after receiving a reachability confirmation message. The default is 30000.
RetransmitTime	The time (in milliseconds) between retransmissions of neighbor solicitation messages to a neighbor when resolving the address or when probing the reachability of a neighbor. The default is 1000.
MulticastAdminStatus	The indication of whether multicasting for IPv6 is enabled (up) or disabled (down) on this interface. The default is false.

Displaying statistics for an IPv6 interface

Use this procedure to export statistics to another application, and to display and graph statistics for an IPv6 interface.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. In the work area, click the **Interfaces** tab.
4. Select an interface from the list.

5. In the work area, select data to export to another application. You can use **Ctrl+click** to select more than one row of data.
6. On the toolbar, you can click **Export** to export the data to another application, or you can click the **Graph** button to open the graph window.
7. If you clicked the **Graph** button, on the tool bar click a graph type to display a graph of the selected IPv6 interface statistics.

Configuring an IPv6 address for a VLAN

Use this procedure to assign site-local or global IPv6 addresses to a VLAN, enabling IPv6 routing for the interface.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. Click the **Addresses** tab.
4. Click **Insert**.
5. In the **IfIndex** box, click **VLAN** , and select a VLAN.
6. Edit the remaining fields.
7. Click **Insert** .
8. Click **Apply** .

Variable definitions

The following table describes the fields of the **Addresses** tab.

Variable	Value
IfIndex	The index value that uniquely identifies the interface to which this entry applies.
Addr	<p>The IPv6 address to which this entry addressing information pertains.</p> <p>! Important:</p> <p>If the IPv6 address exceeds 116 octets, the object identifiers (OIDS) of instances</p>

Variable	Value
	of columns in this row are more than 128 subidentifiers and you cannot use SNMPv1, SNMPv2c, or SNMPv3 to access them.
AddrLen	The prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Type	The type of address: unicast or anycast. The default is unicast.
Origin	A read-only value indicating the origin of the address. The origin of the address is other, manual, dhcp, linklayer, or random.
Status	A read-only value indicating the status of the address, describing whether the address is used for communication. The status is preferred (default), deprecated, invalid, inaccessible, unknown, tentative, or duplicate.

Configuring an IPv6 discovery prefix

The IPv6 discovery prefix determines the source of an IP address or set of IP addresses. The discovery prefix also permits other tables to share the information through a pointer rather than by copying. For example, when the node configures both a unicast and anycast address for a prefix, the `ipAddressPrefix` objects for those addresses point to a single row in the table.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. Click the **Discovery Prefix** tab.
4. Click **Insert**.
5. In the **IfIndex** box, click **Port** or **VLAN**, and select a port number or VLAN.
6. Edit the remaining fields.
7. Click **Insert**.
8. Click **Apply**.

Variable definitions

The following table describes the fields of the **Discovery Prefix** tab.

Variable	Value
IfIndex	A read-only value indicating the unique value to identify an IPv6 interface.
Prefix	Configures the prefix to create an IPv6 address in the IPv6 interface table.
PrefixLen	Configures the mask to create an IPv6 prefix entry as either advertised or suppressed.
VlanId	Specifies the VLAN ID of the IPv6 interface.
UseDefaultVal	Select one of the values to set its value to default value. This is a bitmask field, setting all the bits means that all the options will be reverted to default values.
ValidLife	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.
PreferredLife	Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.
Infinite	Configures the prefix valid lifetime so it never expires. The default is false.
OnLinkFlag	Configures the prefix for use when determining if a node is onlink. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.
AutoFlag	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.
AddressEui	Configures the EUI address. Use an EUI-64 interface ID in the low-order 64-bits of the address when the ID is not specified in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both global and link-local addresses. After you create the entry, you cannot modify this

Variable	Value
	value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used.
NoAdvertise	Select true to not include the prefix in the neighbor advertisement. The default is false.

Configuring route advertisement

Configure router advertisement in IPv6 for neighbor discovery (ND). IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. Click the **Route Advertisement** tab.
4. Edit the fields as required.
5. Click **Apply**.

Variable definitions

The following table describes the fields of the **Route Advertisement** tab.

Variable	Value
IfIndex	A unique value to identify a physical interface or a logical interface (VLAN).
SendAdverts	Indicates whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Select one included value to use the default value, or use all bits to configure all options to their default value.

Variable	Value
MaxInterval	Configure the maximum interval (in seconds) at which the transmission of router advertisements occurs on this interface. This must be no less than 4 seconds and no greater than 1800 seconds. The default is 600.
MinInterval	Configure the minimum interval (in seconds) at which the transmission of router advertisements can occur on this interface. The value must be no less than 3 seconds and no greater than .75 x max-interval. The default is 200.
ReachableTime	The value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this router). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. The default is 30000.
RetransmitTimer	The value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this router). The value configures the amount of time that router waits for the transmission to occur. The default is 1000.
DefaultLifeTime	The value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between rclpv6RouterAdvertMaxInterval and 9000 seconds. A value of zero indicates that the router is not a default router. The default is 3 times the value of rclpv6RouterAdvertMaxInterval or 1800.
CurHopLimit	The default value placed in the current hop limit field in router advertisements sent from this interface. The value must be the current diameter of the Internet. A value of zero in the router advertisement indicates that the advertisement is not specifying a value for CurHopLimit. The value must be the value specified in the IANA Web pages (www.iana.org). The default is 30.
ManagedFlag	If enabled, the ManagedFlag configures the M-bit or the managed address configuration

Variable	Value
	in the router advertisement. The default is false.
OtherConfigFlag	If set to true, then the O-bit (Other stateful configuration) in the router advertisement is set. Reference RFC2461 Section 6.2.1. The default value is false.
DadNSNum	The number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	The value placed in MTU options sent by the router on this interface. A value of zero indicates that the router sends no MTU options.

Creating IPv6 static routes

Use this procedure to configure static routes to destination IPv6 address prefixes.


Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. Click the **Static Routes** tab.
4. Click **Insert**.
5. In the **Dest** box, type the IPv6 address.
6. In the **PrefixLength** box, type the length of the prefix for the IPv6 address.
7. In the **NextHop** box, type the IPv6 address of the router through which the specified route is accessible.
8. In the **IfIndex** box, click **VLAN**, or **Tunnel** and select an option.
9. In the **Cost** box, type a number for the distance.
10. Select the **Enable** check box.
11. In the **Preference** box, type the route preference.
12. Click **Insert**.

The new route appears in the **Static Routes** tab.

Variable definitions

The following table describes the fields of the **Static Routes** tab.

Variable	Value
Dest	Configures the IPv6 destination network address. The prefix value must match the PrefixLength.
PrefixLength	Configures the number of leading one bits that form the mask as a logical value. The prefix value must match the value in the Dest box. The range is 0–128.
NextHop	Configures the next hop IPv6 address.
IfIndex	Select the required VLAN, or tunnel.
Cost	Configures the cost or distance ratio to reach the destination for this node. The range is 1–65535. The default value is 1.
Enable	Configures whether the configured static route is available on the port. The default is enable.  Important: If a static route is disabled, you must enable it before you can add the route to the system routing table.
Status	Indicates the current status of this entry.
Preference	Configures the routing preference of the destination IPv6 address. The range is 1-255. The default value is 5.

Configuring the neighbor cache

Neighbor cache in IPv6 is similar to the IPv4 Address Resolution Protocol (ARP) table. The neighbor cache is a set of entries for individual neighbors to which traffic was sent recently. You make entries on the neighbor on-link unicast IP address, including information such as the link-layer address. A neighbor cache entry contains information used by the Neighbor Unreachability Detection algorithm, including the reachability state, the number of unanswered probes, and the time the next Neighbor Unreachability Detection event is scheduled.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. Click the **Neighbors** tab.
4. Click **Insert**.
5. In the **IfIndex** box, click **Port in VLAN**, and select a VLAN.
6. Edit the remaining fields.
7. Click **Insert**.
8. Click **Apply**.

Variable definitions

The following table describes the fields of the **Neighbors** tab.

Variable	Value
IfIndex	A unique value to identify a logical interface (VLAN). For the VLAN, the value is the ifindex of the vlan.
NetAddress	The IP address corresponding to the media-dependent physical address.
PhyAddress	The media-dependent physical address. The range is 0–65535. For Ethernet, this is a MAC address.
Interface	Either a physical port ID or the MLT port ID. This entry is associated either with a port or with the MLT in a VLAN.
LastUpdated	The value of sysUpTime at the time this entry was last updated. If this entry was updated prior to the last reinitialization of the local network management subsystem, this object contains a zero value.
Type	The mapping type is as follows: <ul style="list-style-type: none"> • Dynamic type: indicates that the IP address to the physical address mapping was dynamically resolved using, for

Variable	Value
	<p>example, IPv4 ARP or the IPv6 Neighbor Discovery Protocol.</p> <ul style="list-style-type: none"> • Static type: indicates that the mapping was statically configured. • Local type: indicates that the mapping is provided for the interface address. <p>The default is static.</p>
State	<p>The Neighbor Unreachability Detection state for the interface when the address mapping in this entry is used. If Neighbor Unreachability Detection is not in use (for example, for IPv4), this object is always unknown. Options include the following:</p> <ul style="list-style-type: none"> • reachable: confirmed reachability • stale: unconfirmed reachability • delay: waiting for reachability confirmation before entering the probe state • probe: actively probing • invalid: an invalidated mapping • unknown: state cannot be determined • incomplete: address resolution is being performed

Configuring the IPv4 remote access list using EDM

Use this procedure to configure a list of IPv4 source addresses for which to permit remote access to a switch.

Procedure steps

1. From the navigation tree, double-click **Administration**.
2. In the Administration tree, double-click **Remote Access**.
3. In the work area, click the **Allowed List(IPv4)** tab.
4. To select a source to edit, click the source row.
5. In the source row double-click the cell in the **Allowed Source IP Address** column.
6. In the dialog box, type a value.

7. In the source row double-click the cell in the **Allowed Source Mask** column.
8. In the dialog box, type a value.
9. Click **Apply** .

Use the data in this table to configure to configure a list of IPv4 source addresses for which to permit access to the switch.

Table 12: Variable definitions

Variable	Value
Allowed Source IP Address	Specifies the source IPv4 address to permit remote access to the switch.
Allowed Source Mask	Specifies subnet mask associated with the source IPv4 address to permit remote access to the switch.

Configuring the IPv6 remote access list using EDM

Use this procedure to configure a list of IPv6 source addresses for which to permit remote access to a switch.

Procedure steps

1. From the navigation tree, double-click **Administration**.
2. In the Administration tree, double-click **Remote Access** .
3. In the work area, click the **Allowed List(IPv6)** tab.
4. To select a source to edit, click the source row.
5. In the source row double-click the cell in the **Allowed Source IPv6 Address** column.
6. In the dialog box, type a value.
7. In the source row double-click the cell in the **Allowed Prefix Length** column.
8. In the dialog box, type a value.
9. Click **Apply** .

Use the data in this table to configure to configure a list of IPv6 source addresses for which to permit access to the switch .

Table 13: Variable definitions

Variable	Value
Allowed Source IPv6 Address	Specifies the source IPv6 address to permit remote access to the switch.

Variable	Value
Allowed Prefix Length	Specifies prefix length for the source IPv6 address to permit remote access to the switch. Values range from 0 to 128.

Graphing IPv6 interface ICMP statistics using EDM

Use this procedure to display and graph IPv6 interface ICMP statistics.

Procedure

1. In the navigation tree, double-click **IPv6** to open the IPv6 navigation tree.
 2. In the IPv6 navigation tree, click **IPv6**.
 3. In the work area, click the **ICMP Stats** tab.
 4. To clear the interface statistics counters, click **Clear Counters**.
 5. Click the arrow on the **Poll Interval:** box.
 6. Select a value from the list.
 7. To select data to graph, click a data row under a column heading.
 8. Click **Line Chart**, **Area Chart**, **Bar Chart**, or **Pie Chart**.
-

ICMP Stats tab field descriptions

The following table describes the fields on the ICMP Stats tab.

Name	Description
InMsgs	Number of ICMP messages received.
InErrors	Number of ICMP error messages received
OutMsgs	Number of ICMP messages sent.
OutErrors	Number of ICMP error messages sent.
Poll Interval	Sets polling interval. Value: 2 to 60 s.

Viewing ICMP message statistics using EDM

Use this procedure to view the IPv6 interface ICMP message statistics using EDM

Procedure

1. In the navigation tree, double-click **IPv6** to open the IPv6 navigation tree.
 2. In the IPv6 navigation tree, click **IPv6**.
 3. In the work area, click the **ICMP Msg Stats** tab.
 4. Click **Refresh** to update the ICMP message statistics.
-

Variable definitions

The following table describes the fields on the ICMP Msg Stats tab.

Name	Description
Type	Type of packet received or sent.
InPkts	Number of packets received.
OutPkts	Number of packets sent.

Viewing global IPv6 TCP properties using EDM

Use this procedure to view IPv6 TCP properties for the switch.

Procedure

1. In the navigation tree, double-click **IPv6** to open the IPv6 navigation tree.
 2. In the IPv6 navigation tree, click **TCP/UDP**.
 3. In the work area, click the **TCP Globals** tab.
 4. Click **Refresh** to update the information.
-

TCP Globals tab field descriptions

The following table describes the fields on the TCP Globals tab.

Name	Description
RtoAlgorithm	Algorithm identifier.
RtoMin	Minimum value in milliseconds.
RtoMax	Maximum value in milliseconds.
MaxConn	Maximum number of connections.

Viewing IPv6 TCP connections using EDM

Use this procedure to view IPv6 TCP connections using EDM.

Procedure

1. In the navigation tree, double-click **IPv6** to open the IPv6 navigation tree.
 2. In the IPv6 navigation tree, click **TCP/UDP**.
 3. In the work area, click the **TCP Connections** tab.
 4. Click **Refresh** to update the information.
-

TCP Connections tab field descriptions

The following table describes the fields on the TCP connections tab.

Name	Description
LocalAddressType	Local address type
LocalAddress	Local address
LocalAddress Port	Local address port
LocalPort	Local port IP
RemAddress Type	Remote address type
RemAddress	Remote address
RemPort	Remote port IP

Name	Description
State	State <ul style="list-style-type: none">• Enabled• Disabled

Viewing IPv6 TCP listeners using EDM

Use this procedure to view IPv6 TCP listeners using EDM.

Procedure

1. In the navigation tree, double-click **IPv6** to open the IPv6 navigation tree.
 2. In the IPv6 navigation tree, click **TCP/UDP**.
 3. In the work area, click the **TCP Listeners** tab.
 4. Click **Refresh** to update the information.
-

TCP Listeners tab field descriptions

The following table describes the fields on the TCP Listeners tab.

Name	Description
LocalAddressType	Local address type
LocalAddress	Local address
Local Port	Local port

Viewing IPv6 UDP endpoints using EDM

Use this procedure to view IPv6 UDP endpoints using EDM

Procedure

1. In the navigation tree, double-click **IPv6** to open the IPv6 navigation tree.
2. In the IPv6 navigation tree, click **TCP/UDP**.

3. In the work area, click the **UDP Endpoints** tab.
 4. Click **Refresh** to update the information.
-

UDP Endpoints tab field descriptions

The following table describes the fields on the UDP Endpoints tab.

Name	Description
LocalAddressType	Local address
LocalAddress	Local address port
Local Port	Local port IP
RemoteAddressType	Remote address type
RemoteAddress	Remote address
RemotePort	Remote port IP
Instance	Indicates the instance.
Process	Indicates the process.

Chapter 42: IPv6 DHCP Relay configuration using ACLI

This chapter describes how to use the ACLI to configure IPv6 DHCP Relay.

IPv6 DHCP Relay configuration procedures

To configure IPv6 DHCP Relay, perform the following steps:

1. Specify the local relay agent and remote server.
2. Enable IPv6 DHCP Relay on the VLAN.

Specifying a local DHCP relay agent and remote DHCP server

Use this procedure to specify a 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.

Procedure steps

To configure a VLAN as a DHCP relay agent, enter the following from the Global Configuration mode:

```
[no] ipv6 dhcp-relay fwd-path <ipv6-relay-agent> <DHCP-server>
[enable]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified DHCP forwarding path.
<ipv6-relay-agent>	Specifies the IPv6 address of the VLAN that serves as the local DHCP relay agent.
<DHCP-server>	Specifies the IPv6 address of the remote DHCP server to which DHCP packets are to be relayed.
[enable]	Enables the specified DHCP relay forwarding path.

You can also specify a VLAN as the DHCP relay agent on the forwarding path to a remote DHCP server from the VLAN Interface Configuration mode:

From the VLAN Interface Configuration mode, enter the following command:

```
[no] ipv6 dhcp-relay fwd-path <DHCP-server> [enable]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified DHCP forwarding path.
<DHCP-server>	Specifies the IPv6 address of the remote DHCP server to which DHCP packets are to be relayed.
[enable]	Enables the specified DHCP relay forwarding path.

Displaying the DHCP relay configuration

Use this procedure to display the current DHCP relay agent configuration.

Procedure steps

To display the DHCP relay configuration, enter the following from the User EXEC command mode:

```
show ipv6 dhcp-relay fwd-path
```

Job aid

The following table shows the field descriptions for the `show ipv6 dhcp-relay fwd-path` command.

Field	Description
INTERFACE	Specifies the interface IPv6 address of the DHCP relay agent.
SERVER	Specifies the IPv6 address of the DHCP server.
ENABLE	Specifies whether DHCP is enabled.

Configuring DHCP relay status and parameters on a VLAN

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.

Procedure steps

To configure DHCP relay on a VLAN, enter the following from the VLAN Interface Configuration mode:

```
[no] ipv6 dhcp-relay [max-hop <max-hop>] [remote-id]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables DHCP relay on the specified VLAN.
[max-hop <max-hop>]	Configures the max hop count, from 1-32.
[remote-id]	Enables remote ID.

Displaying the DHCP relay configuration for a VLAN

Use this procedure to display the current DHCP relay parameters configured for a VLAN.

Procedure steps

To display the DHCP relay VLAN parameters, enter the following from the Privileged EXEC command mode:

```
show ipv6 dhcp-relay interface [vlan <vid>]
```

Variable definitions

The following table describes the 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 ipv6 dhcp-relay` command.

Field	Description
VLAN ID	Indicates the VLAN ID
IfIndex	Indicates the VLAN interface index.
MAX HOP	Indicates the maximum hop value.

Field	Description
DHCP-RELAY	Indicates whether DHCP relay is enabled on the VLAN.
REMOTE ID	Indicates whether Remote ID is enabled on the VLAN. This release does not support Remote ID.

Displaying DHCP relay counters

Use this procedure to display the current DHCP relay counters. This includes the number of requests and the number of replies.

Procedure steps

To display the DHCP relay counters, enter the following from the PrivEXEC command mode:

```
show ipv6 dhcp-relay counters
```

Job aid

The following table shows the field descriptions for the `show ipv6 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.

Chapter 43: IPv6 Tunnel configuration using ACLI

This chapter describes how to use the ACLI to configure IPv6 tunnels.

IPv6 tunnel configuration procedures

To configure IPv6 tunnels, perform the following steps:

1. Configure the tunnel at the source and destination switch.
2. Configure static routes at the source and destination switch.
3. Optionally, configure the tunnel hop limit.

Configuring manual IPv6-in-IPv4 tunnels using the ACLI

Create an IPv6-in-IPv4 tunnel to transfer traffic between IPv6 devices across an IPv4 network. To configure a manual tunnel, you must define both the local and destination IPv4 addresses and configure a static route to route traffic to the IPv6 destination. You must also configure the tunnel at both the source and destination nodes. The maximum number of tunnels is four.

Procedure steps

1. To configure the tunnel, at the source and destination nodes, enter the following in Global Configuration mode:

```
[no] ipv6 tunnel <tunnel id> source <A.B.C.D> address <ipv6 address/prefix-len> destination <A.B.C.D>
```
2. To configure the hop limit, enter the following in Global Configuration mode:

```
ipv6 tunnel <tunnel id> hop-limit <value>
```
3. To utilize the manual tunnels, you must add a static route for the remote IPv6 address. To configure the static route, enter the following in Global Configuration mode:

```
ipv6 route <dest-ipv6address/prefix> tunnel <tunnel-id>
[enable] [cost <1-65535>] [preference <1-255>]
```

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified tunnel.
<tunnel id>	Specifies the ID number of the tunnel in the range of 1-2147483647.
source <A.B.C.D>	Configures the IPv4 source address for the local tunnel.
address <ipv6-address/prefix-len>	Configures the IPv6 source address for the local tunnel in IPv6/prefix-length format.
destination <A.B.C.D>	Specifies the remote IPv4 address for the tunnel destination.
hop-limit <value>	Configures the maximum number of hops that a packet can make before it is dropped. Value is in the range 0–255. To set this option to the default value, use the default operator with the command. The default value is 64.
<dest-ipv6address/prefix>	Specifies the IPv6 address of the remote IPv6 tunnel destination.
enable	Enables the route.
cost <1-65535>	Specifies the metric of the route in the range of 1–65535.
preference <1-255>	Specifies the route preference in the range of 1–255. The default value is 0.

Displaying manual tunnel configuration

Use the following procedure to display the manual tunnel configuration.

Procedure steps

To display the configured tunnel, enter the following in PrivExec mode:


```
show ipv6 tunnel [<tunnel-id>]
```


Chapter 44: IPv6 DHCP Relay configuration using Enterprise Device Manager

This chapter describes how to configure IPv6 DHCP Relay using EDM.

IPv6 DHCP Relay configuration procedures

To configure IPv6 DHCP Relay, perform the following steps:

1. Specify the local relay agent and remote server.
2. Enable IPv6 DHCP Relay on the VLAN.

Configuring the DHCP relay forwarding path

Configure forwarding policies to indicate the relay agent and the DHCP server to which packets are forwarded.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **DHCP Relay**.
3. Click the **Forward Path** tab.
4. Click **Insert**.
5. In the **AgentAddr** box, type the agent address.
6. In the **ServerAddr** box, type the server address.
7. Click **Enabled** to enable DHCP Relay. You can enable or disable each agent server forwarding path. The default is enabled.
8. Click **Insert**.

Variable definitions

The following table describes the fields of the **Forward Path** tab.

Variable	Value
AgentAddr	The IP address of the relay agent on which the DHCP request packets are received for forwarding. This address is the IP address of a VLAN for which forwarding is enabled.
ServerAddr	This parameter is the IP address of the DHCP server. The request is unicast to the server address.
Enabled	Enables DHCP relay on the routing switch.

Configuring DHCP relay interface parameters

Configure the DHCP relay behavior on the interface.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **DHCP Relay**.
3. In the Interface tab, click **Insert**.
4. Enter the appropriate values.
5. Click **Insert**.

Variable definitions

The following table describes the fields of the **Interface** tab.

Variable	Value
IfIndex	A read-only value indicating the unique value to identify an IPv6 interface.

Variable	Value
MaxHop	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server.
RemoteldEnabled	Enables or disables remote ID.

Displaying DHCP Relay statistics

Display DHCP Relay statistics to monitor network performance.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **DHCP Relay**.
3. In the **Interface** tab, select an interface and click the **Stats**.

Variable definitions

The following table describes the fields of the **Stats** tab.

Variable	Value
NumRequests	The count of request messages.
NumReplies	The count of reply messages.

Chapter 45: IPv6 Tunnel configuration using Enterprise Device Manager

This chapter describes how to configure IPv6 tunnels using EDM.

IPv6 tunnel configuration procedures

To configure IPv6 tunnels, perform the following steps:

1. Configure the tunnel at the source and destination switch.
2. Configure static routes at the source and destination switch.
3. Optionally, configure the tunnel hop limit.

Configuring a tunnel for IPv6 VLANs to communicate through an IPv4 network

Use the following procedure to configure a tunnel for IPv6 VLANs to communicate through an IPv4 network. Manual tunnels are point-to-point, so you configure both source and destination addresses. You must configure both IPv6 and IPv4 addresses for both source and destination devices. The IPv6 addresses must represent the same network, for example 6666::1/96 and 6666::2/96.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **Tunnel**.
3. Click **Insert**.
4. In the **LocalAddress** box, click the button and select the IPv4 address for the local VLAN.

5. In the **RemoteAddress** box, type the IPv4 address for the remote destination VLAN.
6. In the **EncapsMethod** area, select **manual**.
7. In the **ID** box, type a number to represent the tunnel.
8. In the **IPv6AddressAddr** box, type the IPv6 address assigned to the tunnel.
9. In the **IPv6AddressPrefixLength** box, type the number of bits to advertise in the IPv6 address.
10. Click **Insert**.

After you create the tunnel, the **Local Address** tab displays the IPv4 addresses associated with the tunnel.

11. To view the configured IPv6 Address for the tunnel, click **IPv6 Address**.

Variable definitions

The following table describes the fields of the **Tunnel Config** tab.

Variable	Value
Address Type	Displays the address type for the tunnel: IPv4 for IPv6 packets encapsulated in IPv4.
LocalAddress	Identifies the local endpoint address of the tunnel.
RemoteAddress	Identifies the remote endpoint of the tunnel.
EncapsMethod	Displays the tunnel mode: manual for manually configured tunnels.
ID	Identifies the tunnel number.
IfIndex	Displays a unique value that identifies the tunnel interface internally. The value is derived from the tunnel ID.

Viewing the local IPv6 address associated with a tunnel

Use this procedure to view the local IPv6 address associated with a pre-configured tunnel.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **Tunnel**.
3. Select the desired tunnel.
4. To view the configured IPv6 Address for the tunnel, click **IPv6 Address**.

Variable definitions

The following table describes the fields of the **IPv6 Address** tab.

Variable	Value
IfIndex	The index value that uniquely identifies the interface to which this entry applies.
Addr	<p>The IPv6 address to which this entry addressing information pertains.</p> <p>* Note:</p> <p>If the IPv6 address exceeds 116 octets, the object identifiers (OIDs) of instances of columns in this row are more than 128 subidentifiers and you cannot use SNMPv1, SNMPv2c, or SNMPv3 to access them.</p>
AddrLen	The prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Type	The type of address: unicast or anycast. The default is unicast.
Origin	<p>A read-only value indicating the origin of the address.</p> <p>The origin of the address is one of the following:</p> <ul style="list-style-type: none"> • other • manual • dhcp • linklayer • random

Variable	Value
Status	<p>A read-only value indicating the status of the address, describing whether the address is used for communication. The status is one of the following:</p> <ul style="list-style-type: none"> • preferred (the default) • deprecated • invalid • inaccessible • unknown • tentative • duplicate

Modifying tunnel hop limits

Modify tunnel hop limits to update hop limit values on previously configured tunnels.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
 2. Double-click **Tunnel**.
 3. Click the **Tunnel Interfaces** tab.
- OR**
- Click the **Tunnel Interface** button.
4. In the row for the tunnel to configure, double-click the **HopLimit** column to modify the displayed information as required.
 5. Click **Apply**.

Variable definitions

The following table describes the fields of the **Tunnel Interface** tab.

Variable	Value
Index	Identifies the tunnel interface internally. The value is derived from the tunnel ID.
EncapsMethod	Displays the tunnel mode: IPv6 for manually configured tunnels and 6to4 for automatically configured tunnels.
HopLimit	Configures the maximum number of hops in the tunnel. The default value is 255.
Security	Indicates the type of security on the tunnel interface.
TOS	<p>Displays the method used to configure the high 6 bits (the differentiated services codepoint) of the IPv4 type of service (TOS) or IPv6 traffic class in the outer IP header. A value of —1 indicates that the bits are copied from the payload header.</p> <p>A value of —2 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB module.</p> <p>A value from 0 to 63 indicates that the bit field is configured to the indicated value.</p>
FlowLabel	<p>Displays the method used to set the IPv6 Flow Label value. This object need not be present in rows where tunnelIfAddressType indicates that the tunnel is not over IPv6.</p> <p>A value of —1 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB.</p> <p>Any other value indicates that the Flow Label field is configured to the indicated value.</p>
AddressType	<p>Displays one of the following values:</p> <ul style="list-style-type: none"> • Manual — a manually configured tunnel • 6to4 for autoconfigured tunnels.
LocalNetAddress	Identifies the local endpoint address of the tunnel.
RemoteNetAddress	Identifies the remote endpoint of the tunnel.
EncapsLimit	<p>Displays the address of the local endpoint of the tunnel (that is, the source address used in the outer IP header). If the address is unknown, the value is 0.0.0.0 for IPv4 or :: for IPv6.</p> <p>The tunnelIfAddressType displays the object type.</p>

Creating IPv6 static routes

Use this procedure to configure static routes to destination IPv6 address prefixes.

Procedure steps

1. From the navigation tree, double-click **IPv6**.
2. Double-click **IPv6**.
3. Click the **Static Routes** tab.
4. Click **Insert**.
5. In the **Dest** box, type the IPv6 address of the tunnel destination.
6. In the **PrefixLength** box, type the length of the prefix for the IPv6 address.
7. In the **NextHop** box, type the IPv6 address of the local IPv6 address through which the specified tunnel is accessible.
8. In the **IfIndex** box, click **Tunnel** and select the previously configured Tunnel ID.
9. In the **Cost** box, type a number for the distance.
10. Select the **Enable** check box.
11. In the **Preference** box, type the route preference.
12. Click **Insert**.

The new route appears in the **Static Routes** tab.

The following table describes the fields of the **Static Routes** tab.

Table 14: Variable definitions

Variable	Value
Dest	Configures the IPv6 destination network address. The prefix value must match the PrefixLength.
PrefixLength	Configures the number of leading one bits that form the mask as a logical value. The prefix value must match the value in the Dest box. The range is 0–128.
NextHop	Configures the next hop IPv6 address.
IfIndex	Select the required VLAN, or tunnel.
Cost	Configures the cost or distance ratio to reach the destination for this node. The range is 1–65535. The default value is 1.

Variable	Value
Enable	Configures whether the configured static route is available on the port. The default is enable. ! Important: If a static route is disabled, you must enable it before you can add the route to the system routing table.
Status	Indicates the current status of this entry.
Preference	Configures the routing preference of the destination IPv6 address. The range is 1-255. The default value is 5.

