Configuring IP Routing Protocols for Avaya Virtual Services Platform 9000

Release 4.0
NN46250-505
Issue 06.01
December 2014

© 2014 Avaya Inc.

All Rights Reserved.

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

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

Licenses

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

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. “Software” means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users. “Instance” means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine (“VM”) or similar deployment.

License types

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

Heritage Nortel Software

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

Copyright

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

Third Party Components

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

Preventing Toll Fraud

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

result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

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

Downloading Documentation

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

Contact Avaya Support

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

Contents

Chapter 1: Introduction 9
Purpose 9
Related resources 9
Documentation 9
Training 10
Viewing Avaya Mentor videos 10
Support 10
Searching a documentation collection 11
Chapter 2: New in this release 12
Features 12
Other changes 14
Chapter 3: IP routing operations fundamentals 15
IP addressing 15
IP address for the management port 19
Loopback 19
Static routes 20
Black hole static routes 21
VLANs and routing 21
Equal Cost Multipath 22
Alternate route 23
Route filtering and IP policies 24
Prefix list 27
Route policy definition 28
Individual port routing control 32
Reverse path checking 32
Address Resolution Protocol 34
Reverse Address Resolution Protocol 36
DHCP option 82 36
UDP broadcast forwarding 39
Virtual Router Redundancy Protocol 40
RSMLT 45
Chapter 4: ARP configuration using ACLI 50
Enabling ARP on a port or a VLAN 50
Enabling ARP proxy 51
Configuring ARP loop detection 51
Showing ARP information 52
Configuring IP ARP static entries 54
Clearing ARP entries 55
Showing ARP table information 56

Chapter 5: ARP configuration using Enterprise Device Manager 60
Enabling or disabling ARP on the brouter port or a VRF instance 60
Enabling or disabling ARP on a VLAN or a VRF instance 61
Viewing and managing ARP 62
Creating static ARP entries 63
Configuring ARP proxy 63
Chapter 6: DHCP and UDP configuration using ACLI 65
Configuring DHCP parameters globally 65
Showing DHCP relay information 67
Configuring DHCP option 82 68
Configuring DHCP relay on a port or VLAN 70
Configuring UDP broadcast forwarding 72
Configuring UDP protocols 72
Configuring a UDP port forward entry 73
Configuring the UDP port forwarding list 74
Showing UDP forward information 77
Chapter 7: DHCP and UDP configuration using Enterprise Device Manager 79
Configuring DHCP on a brouter port or a VRF instance 79
Configuring BootP/DHCP on a VLAN or VRF instance 81
Configuring DHCP relay 82
Viewing DHCP relay configuration information 83
Managing UDP forwarding protocols 84
Managing UDP forwarding 85
Creating the forwarding profile 86
Managing the broadcast interface 87
Viewing UDP endpoint information 88
Chapter 8: IP policy configuration using ACLI 90
Configuring prefix lists 90
Configuring IP route policies 92
Job aid 98
Configuring a policy to accept external routes from a router 98
Applying OSPF accept policy changes 100
Configuring inter-VRF redistribution policies 101
Chapter 9: IP policy configuration using Enterprise Device Manager 104
Configuring a prefix list 104
Configuring a route policy 105
Applying a route policy 109
Viewing IP routes 110
Configuring an OSPF accept policy 112
Configuring inbound/outbound filtering policies on a RIP interface 113
Deleting inbound/outbound filtering policies on a RIP interface 114
Chapter 10: IP routing configuration using ACLI 116

Enabling routing globally or on a VRF instance..... 116
Enabling routing on a port..... 117
Deleting a dynamically learned route..... 118
Configuring IP route preferences..... 119
Flushing routing tables by VLAN or port..... 121
Configuring an IP address for the management port..... 122
Configuring a virtual IP address for the management port..... 123
Assigning an IP address to a port..... 124
Assigning an IP address to a VLAN..... 125
Viewing IP addresses for all router interfaces..... 126
Configuring IP routing globally or for a VRF..... 127
Configuring static routes..... 132
Configuring a black hole static route..... 134
Configuring a default static route..... 135
Enabling ICMP Router Discovery globally..... 137
Configuring Router Discovery on a port or VLAN..... 138
Configuring a CLIP interface..... 140
Chapter 11: IP routing configuration using Enterprise Device Manager..... 142
Enabling routing for a router or a VRF instance..... 142
Enabling or disabling routing on a port..... 142
Deleting a dynamically-learned route..... 143
Configuring IP route preferences..... 145
Flushing routing tables by VLAN..... 146
Flushing routing tables by port..... 147
Assigning an IP address to a port..... 147
Assigning an IP address to a VLAN..... 149
Viewing IP addresses for all router interfaces..... 149
Configuring IP routing features globally..... 150
Configuring ECMP globally..... 153
Enabling alternative routes globally..... 153
Configuring static routes..... 154
Deleting a static route..... 156
Configuring a default static route..... 156
Configuring a black hole static route..... 157
Configuring ICMP Router Discovery globally..... 157
Configuring the ICMP Router Discovery table..... 158
Configuring ICMP Router Discovery for a port..... 159
Configuring ICMP Router Discovery on a VLAN..... 160
Configuring a CLIP interface..... 162
Enabling OSPF on a CLIP interface..... 162
Enabling PIM on a CLIP interface..... 163
Viewing TCP global information..... 164
Viewing TCP connections information..... 165

Viewing TCP listeners information ... 166
Chapter 12: RSMLT configuration using ACLI ... 168
Configuring RSMLT on a VLAN ... 168
Showing IP RSMLT information ... 169
Configuring RSMLT edge support ... 171
Chapter 13: RSMLT configuration using Enterprise Device Manager ... 173
Configuring RSMLT on a VLAN ... 173
Viewing and editing RSMLT local information ... 174
Viewing RSMLT peer information ... 175
Enabling RSMLT Edge support ... 176
Viewing RSMLT edge support information ... 177
Chapter 14: VRRP configuration using ACLI ... 178
Configuring VRRP on a port or a VLAN ... 179
Showing VRRP port or VLAN information ... 182
Showing extended VLAN VRRP ... 185
Showing VRRP interface information ... 187
Configuring VRRP notification control ... 189
Enabling ping to a virtual IP address ... 190
Chapter 15: VRRP configuration using EDM ... 192
Enabling VRRP global variables ... 193
Configuring VRRP for the interface ... 194
Configuring VRRP on a port or a VRF instance ... 196
Configuring VRRP on a VLAN (or brouter port) or a VRF instance ... 198
Configuring Fast Advertisement Interval on a port or a VRF instance ... 201
Configuring Fast Advertisement Interval on a VLAN or a VRF instance ... 201
Chapter 16: VRF Lite fundamentals ... 203
Overview ... 203
VRF Lite capability and functionality ... 204
VRF Lite and inter-VRF route redistribution ... 206
VRF Lite requirements ... 207
Ethernet modules and VRF Lite management ... 207
Management VRF ... 208
VRF Lite configuration rules ... 209
Virtualized protocols ... 211
Chapter 17: VRF Lite configuration using ACLI ... 212
Creating a VRF instance ... 213
Associating a VLAN or port with a VRF instance ... 215
Configuring a static route on the management router ... 217
Creating an IP VPN instance on a VRF ... 218
Chapter 18: VRF Lite configuration using Enterprise Device Manager ... 221
Configuring a VRF instance ... 221
Configuring interVRF route redistribution policies ... 222

Viewing brouter port and VRF associations..... 223
Viewing global VRF status information..... 224
Viewing VRF instance statistics and status information..... 225
Viewing VRF statistics for a VRF..... 226
Selecting and launching a VRF context view 226
Creating an IP VPN instance on a VRF..... 227
Glossary..... 229

Chapter 1: Introduction

Purpose

You can use *Configuration — IP Routing* to configure general routing operations on the Avaya Virtual Services Platform 9000.

Operations include:

- Address Resolution Protocol (ARP)
- reverse ARP
- TCP and UDP
- Dynamic Host Configuration Protocol Relay (DHCP) Relay
- Virtual Router Redundancy Protocol (VRRP)
- VRF-Lite
- circuitless IP (CLIP) interfaces
- static routes
- Point-to-Point Protocol over Ethernet
- Equal Cost Multipath (ECMP)
- Routed Split MultiLink Trunking (RSMLT)
- routing policies

Related resources

Documentation

See *Documentation Reference for Avaya Virtual Services Platform 9000,* NN46250-100 for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the website at http://avaya-learning.com/.

Course code	Course title
4D00010E	Knowledge Access: ACIS - Avaya ERS 8000 and VSP 9000 Implementation
5D00040E	Knowledge Access: ACSS - Avaya VSP 9000 Support

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

Procedure

- To find videos on the Avaya Support website, go to http://support.avaya.com and perform one of the following actions:
  - In **Search**, type `Avaya Mentor Videos` to see a list of the available videos.
  - In **Search**, type the product name. On the Search Results page, select **Video** in the **Content Type** column on the left.
- To find the Avaya Mentor videos on YouTube, go to www.youtube.com/AvayaMentor and perform one of the following actions:
  - Enter a key word or key words in the Search Channel to search for a specific product or topic.
  - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the website.

    Note:

    Videos are not available for all products.

Support

Go to the Avaya Support website at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes,

downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Searching a documentation collection

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

Before you begin

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

Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named <*product_name_release*>.pdx.
3. In the Search dialog box, select the option **In the index named <*product_name_release*>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
   - Whole Words Only
   - Case-Sensitive
   - Include Bookmarks
   - Include Comments
6. Click **Search**.

   The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Chapter 2: New in this release

The following sections detail what is new in *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505 for Release 4.0.

Features

See the following sections for information about feature-related changes for Release 4.0.

ARP tab

Release 4.0 updates the **DestIfIndex** attribute to allow you to configure the value in order to create a static ARP entry using SNMP. For more information, see: Viewing and managing ARP on page 62.

Intermediate-System-to-Intermediate-System (IS-IS) accept policies

Release 4.0 adds Intermediate-System-to-Intermediate-System (IS-IS) accept policies. You can use IS-IS accept policies with Layer 3 VSNs or IP Shortcuts to filter incoming IS-IS routes or accept IS-IS routes from other service identifiers (I-SID) over the SPBM cloud. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table. IS-IS policies can also use either a service instance identifier (I-SID) or an I-SID list to filter incoming traffic.

For information on how to configure IS-IS accept policy filters, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

For more information on how to configure route policies to use with IS-IS accept policy filters, see:

- Route filtering and IP policies on page 24.
- Route policy definition on page 28.
- Configuring IP route policies on page 92.
- Configuring a route policy on page 105.
- Applying a route policy on page 109.

IP Globals tab

Release 4.0 removes the **ArpThreshold** and **ArpMcastMacFlooding** field descriptions from the IP **Globals** tab. For more information, see:

- Globals field descriptions on page 151

IP Routes tab

The document is updated to remove the **AltSequence** option from the IP **Routes** tab and add the **Layer3VirtualInterfaceType** and **Layer3VirtualInterface** options to the IP **Routes** tab. For more information, see:

- Routes field descriptions on page 110.
- Routes field descriptions on page 143.

Loop detection

The document is updated to include the **`default`** option for the **`default loop-detect [action] [arp-detect]`** command and the **`no`** option for the **`no loop-detect [arp-detect]`** command. For more information, see:

- Configuring ARP loop detection on page 51.

Update to show ip route command output

Release 4.0 updates the output for the **`show ip route [vrf WORD<1-16>]`** command. The NH VRF column now displays as NH VRF/ISID and the NH VRF/ISID column displays the I-SID in the following examples:

- Only for inter-Virtual Services Network (VSN) routes leaked using IS-IS accept policies.
- Only if the I-SID for which the routes are leaked does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays.

If the I-SID is 0, which represents the GlobalRouter, the column displays as GlobalRouter.

The existing IS-IS routes in Shortest Path Bridging (SPB) Layer 3 VSN continue to display as the VRF name of the IP VSN.

For more information on IS-IS accept policies, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

For more information on the **`show ip route [vrf WORD<1-16>]`** command, see:

- Deleting a dynamically learned route on page 118.
- Configuring a static route on the management router on page 217 .

Updates to VRF name variable

Release 4.0 updates the VRF name variable to **`vrf WORD<1-16>`** in ACLI. EDM VRF name options are also updated throughout the document to a range of 1–16.

VRRP updates

Minor updates made to add additional parameters, and update the show outputs for the following VRRP procedures. For more information, see:

- Configuring VRRP on a port or a VLAN on page 179.
- Showing extended VLAN VRRP on page 185.
- Showing VRRP interface information on page 187.

Other changes

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

Document title change

In Release 4.0, the title of this document changed from *Avaya Virtual Services Platform 9000 Configuration — IP Routing*, NN46250-505 to *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Globals tab

Release 4.0 **Globals** tab information is updated to display Ecmp5PathList, Ecmp6PathList, Ecmp7PathList, and Ecmp8PathList. For more information, see Globals field descriptions on page 151.

Static ARP entries

Release 4.0 adds information related to static ARP entries. Virtual Services Platform 9000 supports static ARP entries for NLB multicast and NLB multicast IGMP. Virtual Services Platform 9000 does not support static ARP entries for NLB unicast. For more information, see:

- Configuring IP ARP static entries on page 54.
- Creating static ARP entries on page 63.

Chapter 3: IP routing operations fundamentals

Use the information in this section to understand IP routing.

For more information about Border Gateway Protocol (BGP), see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

For more information about Routing Information Protocol (RIP) and Open Shortest Path First (OSPF), see *Configuring OSPF and RIP on Avaya Virtual Services Platform 9000,* NN46250-506.

IP addressing

An IP version 4 address consists of 32 bits expressed in dotted-decimal format (x.x.x.x). The IP version 4 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 IP address space by address range and mask.

Class	Address range	Mask	Number of addresses
A	1.0.0.0 to 126.0.0.0	255.0.0.0	126
B	128.0.0.0 to 191.0.0.0	255.255.0.0	127 * 255
C	192.0.0.0 to 223.0.0.0	255.255.255.0	31 * 255 * 255
D	224.0.0.0 to 239.0.0.0	—	—

To express an IP address in dotted-decimal notation, you convert each octet of the IP address to a decimal number and separate the numbers by decimal points. For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

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

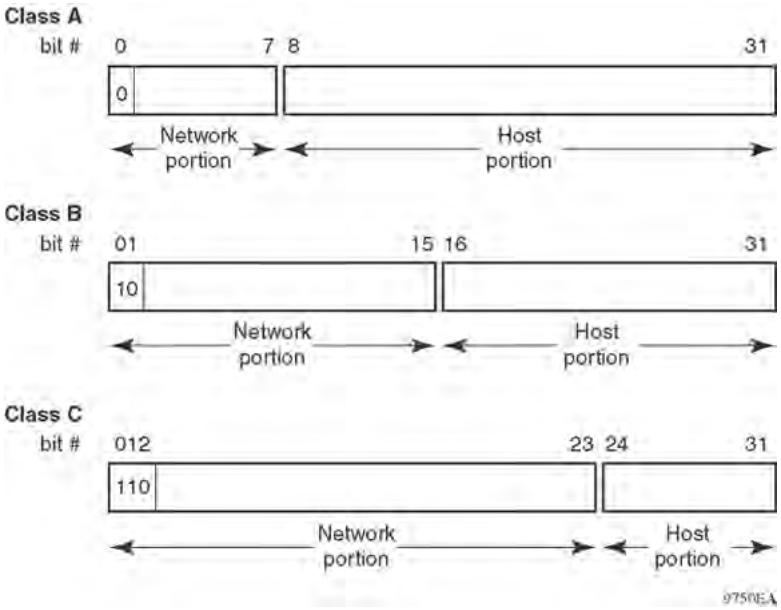


Figure 1: Network and host boundaries in IP address classes

Subnet addressing

Subnetworks (or subnets) extend the IP addressing scheme an organization uses to one with an 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).

You 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 is applied to the IP address and identifies the network and host portions of the address.

The following table illustrates how subnet masks used with class B and class C addresses can create differing numbers of subnets and hosts. This example includes the zero subnet, which is permitted on Virtual Services Platform 9000.

Table 1: Subnet masks for class B and class C IP addresses

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
Class B			
2	255.255.192.0	2	16 382
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

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
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

You use variable-length subnet masking (VLSM) to divide your intranet into pieces that match your requirements. Routing is based on the longest subnet mask or network that matches. Routing Information Protocol version 2 and Open Shortest Path First are routing protocols that support VLSM.

Supernet addressing and CIDR

A supernet, or classless interdomain routing (CIDR) address, is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed. You can use supernetting to address an entire block of class C addresses and avoid using large routing tables to track the addresses.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the class C addresses shown in the following figure. By adding the mask 255.255.128.0 to IP address 192.32.128.0, you aggregate the addresses 192.32.128.0 through 192.32.255.255 and 128 class C addresses use a single routing advertisement. In the bottom half of the following figure, you use 192.32.0.0/17 to aggregate the 128 addresses (192.32.0.0/24 to 192.32.127.0/24).

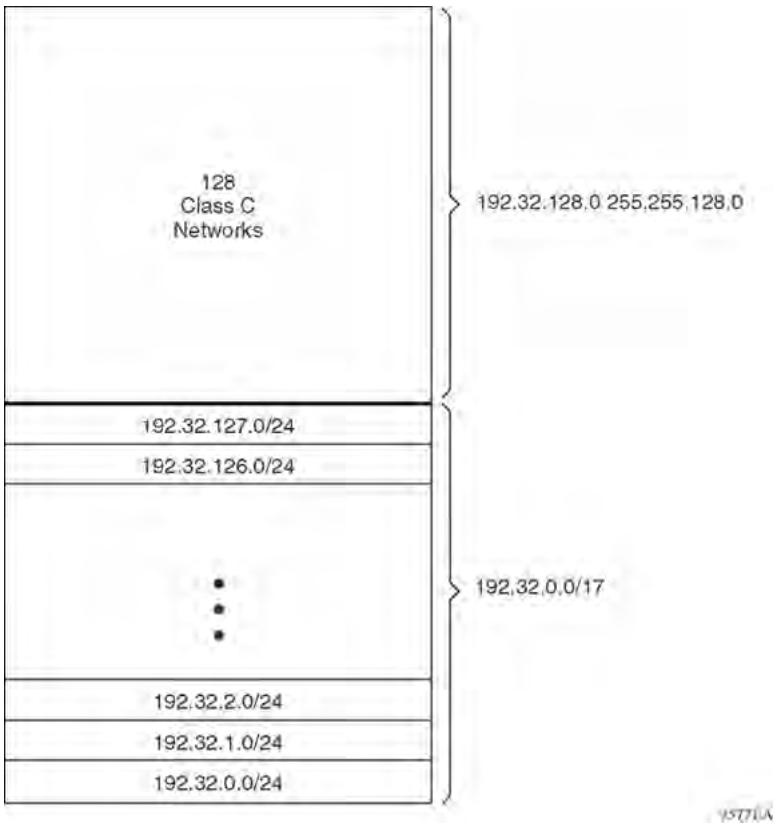


Figure 2: Class C address supernet

Another example is the block of addresses 192.32.0.0 to 192.32.7.0. The supernet address for this block is 11000000 00100000 00000, with the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an address and mask pair:

- The address is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- The mask is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/21.

Although classes prohibit using an address mask with the IP address, you can use CIDR to create networks of various sizes using the address mask. You can also divide the address space using variable-length subnet mask (VLSM); the division is not visible outside your network. With CIDR, the routers outside the network use the addresses.

IP address for the management port

At startup, the system loads the runtime configuration file, which is stored in the internal flash of the Control Processor (CP) module. If the file is present, the system assigns the IP address for the management port from that file.

You can configure an IP address for the management port if one is not in the configuration file. For more information, see Configuring an IP address for the management port on page 122.

Loopback

Circuitless IP (CLIP) is a virtual (or loopback) interface that is not associated with a physical port. You can use the CLIP interface to provide uninterrupted connectivity to your device as long as a path exists to reach the device.

For example, as shown in Figure 3: Routers with IBGP connections on page 19, a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Use an Interior Border Gateway Protocol (IBGP) session between two additional addresses, 195.39.128.1/30 (CLIP 1) and 195.39.281.2/30 (CLIP 2).

CLIP 1 and CLIP 2 represent the virtual CLIP addresses that you configure between R1 and R2. These virtual interfaces are not associated with the physical link or hardware interface, which permits the BGP session to continue as long as a path exists between R1 and R2. An IGP (such as OSPF) routes addresses that correspond to the CLIP addresses. After the routers learn all the CLIP addresses in the AS, the system establishes IBGP and exchanges routes.

You can also use CLIP for PIM-SM, typically, as a Rendezvous Point (RP), or as a source IP address for sending SNMP traps and Syslog messages.

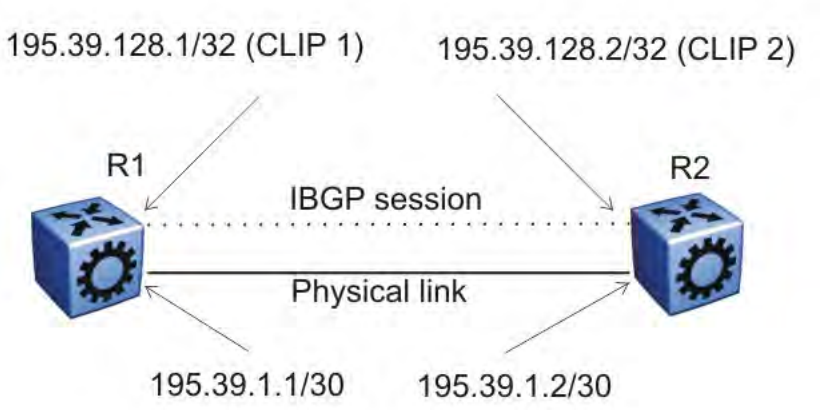


Figure 3: Routers with IBGP connections

The system treats the CLIP interface as an IP interface. The network associated with the CLIP is treated as a local network attached to the device. This route always exists and the circuit is always up because there is no physical attachment.

The system 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.

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

Static routes

A static route is a route to a destination IP address that you manually create.

The Layer 3 redundancy feature supports the creation of static routes to enhance network stability. Use the local next hop option to configure a static route with or without local next hop.

You can configure static routes with a next hop that is not directly connected, but that hop must be reachable. Otherwise, the static route is not enabled.

Layer 3 redundancy supports only address resolution protocol (ARP) and static route. Static ARP must configure the nonlocal next-hop of static routes. No other dynamic routing protocols provide nonlocal next-hop.

You can use a default static route to specify a route to all networks for which no explicit routes exist in the forwarding information base or the routing table. This route has a prefix length of zero (RFC1812). You can configure Avaya Virtual Services Platform 9000 with a route through the IP static routing table.

To create a default static route, you must configure the destination address and subnet mask to 0.0.0.0.

Static route tables

A router uses the system routing table to make forwarding decisions. In the static route table, you can change static routes directly. Although the two tables are separate, the static route table manager entries are automatically reflected in the system routing table if the next-hop address in the static route is reachable, and if the static route is enabled.

The system routing table displays only active static routes with a best preference. A static route is active only if the route is enabled and the next-hop address is reachable (for example, if a valid ARP entry exists for the next hop).

You can enter multiple routes (for example, multiple default routes) that have different costs, and the routing table uses the lowest cost route that is available. However, if you enter multiple next hops for the same route with the same cost, the software 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. If the first route becomes unreachable, the second route (with a different next-hop) is activated with no connectivity loss.

Black hole static routes

A black hole static route is a route with an invalid next hop, and the device drops data packets destined for this network.

While the router aggregates or injects routes to other routers, the router does not have a path to the aggregated destination. In such cases, the result is a black hole and a routing loop. To avoid routing loops, configure a black hole static route to the destination the router is advertising.

You can configure a preference value for a black hole route. However, you must configure that preference value appropriately so that when you want to use the black hole route, it is elected as the best route.

Before you add a black hole static route, perform a check to ensure that no other static route to that identical destination is enabled. If such a route exists, you cannot add the black hole route and an error message appears.

If a you enable a black hole route, you cannot add another static route to that destination. You must first delete or disable the black hole route before you add a regular static route to that destination.

VLANs and routing

When traffic is routed on a virtual local area network (VLAN), an IP address is assigned to the VLAN and is not associated with a particular physical port. Brouter ports are VLANs that route IP packets and bridge nonroutable traffic in a single-port VLAN.

Virtual routing between VLANs

Virtual Services Platform 9000 supports wire-speed IP routing between VLANs. As shown in the following figure, VLAN 1 and VLAN 2 are on the same device, yet for traffic to flow from VLAN 1 to VLAN 2, the traffic must be routed.

When you configure routing on a VLAN, you assign an IP address to the VLAN, which acts as a virtual router interface address for the VLAN (a virtual router interface is not associated with a particular port). You can reach the VLAN IP address through the VLAN ports, and frames are routed from the VLAN through the gateway IP address. Routed traffic is forwarded to another VLAN within the device.

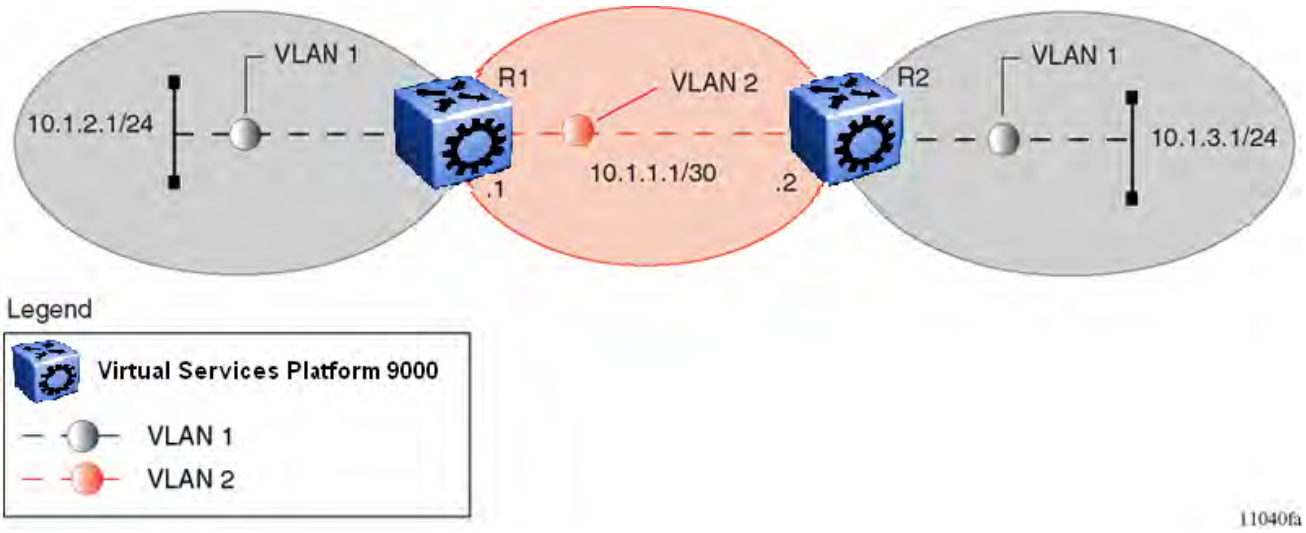


Figure 4: IP routing between VLANs

When Spanning Tree Protocol is enabled in a VLAN, the spanning tree convergence must be stable before the routing protocol begins. This requirement can lead to an additional delay in the IP traffic forwarding.

Because a port can belong to multiple VLANs (some of which are configured for routing on the device and some of which are not), a one-to-one correspondence no longer exists between the physical port and the router interface.

As with an IP address, virtual router interface addresses using Virtual Router Redundancy Protocol (VRRP) are also used for device management. For Simple Network Management Protocol (SNMP) or Telnet management, you can use virtual router interface address to access the device as long as routing is enabled on the VLAN.

Brouter ports

Virtual Services Platform 9000 also supports brouter ports. A brouter port is a single-port VLAN that routes IP packets and bridges all nonroutable traffic. The difference between a brouter port and a standard IP protocol-based VLAN configured to route traffic 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 nonroutable traffic and still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

Because a brouter port is a single-port VLAN, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

Equal Cost Multipath

With Equal Cost Multipath (ECMP), Avaya Virtual Services Platform 9000 can determine up to eight equal-cost paths to the same destination prefix. You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers

more efficiently when sending IP traffic. Equal Cost Multipath is formed using routes from the same source or protocol.

The ECMP feature supports and complements the following protocols and route types:

- OSPF
- Routing Information Protocol (RIP)
- BGP
- Static route
- Default route

Alternate route

Routers can learn several routes to a destination network through several protocols. If you enable the alternate route feature, the Avaya Virtual Services Platform 9000 stores all of these alternate routes sorted in order by network mask, cost, and route preference. The first route on this list is the best route. The hardware uses the first route. The rest of the routes are alternate routes.

To avoid traffic interruption, you can enable alternate routes globally to replace the best route with the next-best route if the best route becomes unavailable. By default, alternate routes are globally enabled. The alternate route concept applies between routing protocols. For example, if an OSPF route becomes unavailable and an alternate RIP route is available, the RIP route is immediately activated without waiting for an update interval to expire.

The internal routing table manager records the route changes for protocols. It maintains separate tables of static (user-configured) and dynamic (protocol-learned) routes and, in the Avaya Virtual Services Platform 9000 software, you can configure preferences that determine the precedence given to one type of route over another.

If a router learns a route with the same network mask and cost values from multiple sources (protocols), the router uses preferences to select the best route to add to the forwarding database. Up to four other routes for each destination are held available as alternative routes.

When you configure a static route on the Avaya Virtual Services Platform 9000, you can specify a preference for the route. To modify the preference for a static route, disable the route before you edit the configuration, and then reenable the route.

Important:

Changing route preferences is a process-intensive operation that can affect system performance and network reachability while you perform route preference procedures. Avaya recommends that if you want to change preferences for static routes or routing protocols, do so when you configure routes or during a maintenance window.

On Virtual Services Platform 9000, default preferences are assigned to all standard routing protocols. You can modify the default preference for a protocol to give it a higher or lower priority than other protocols. When you change the preference for a route, if all best routes remain best

routes, only the local route tables change. However, if changing the protocol preference causes best routes to no longer be best routes, neighboring route tables can be affected.

In addition, you can modify the preference value for dynamic routes through route filtering and IP policies, and this value overrides the global preference for the protocol. You can use alternative mechanisms to change the behavior of specific routes to have a different preference rather than acquiring the global protocol preference. For a static route, you can specify an individual route preference that overrides the global static route preference. The preference value can be between 0 and 255, with 0 reserved for local routes and 255 representing an unreachable route.

The following table shows the default preferences for routing protocols and route types. You can modify the preference value.

Table 2: Routing protocol default preference

Protocol	Default preference
Local	0
Static	5
OSPF intra-area	20
OSPF inter-area	25
Exterior BGP	45
RIP	100
OSPF external type 1	120
OSPF external type 2	125
IBGP	175
Staticv6	5
OSPFv3 intra-area	20
OSPFv3 inter-area	25
OSPFv3 external type 1	120
OSPFv3 external type 2	125

Route filtering and IP policies

When the switch routes IP traffic, you can apply a number of filters to manage, accept, redistribute, and announce policies for unicast routing table information. Filters apply differently to different unicast routing protocols.

The following figure shows how filters apply to BGP, RIP, and OSPF protocols.

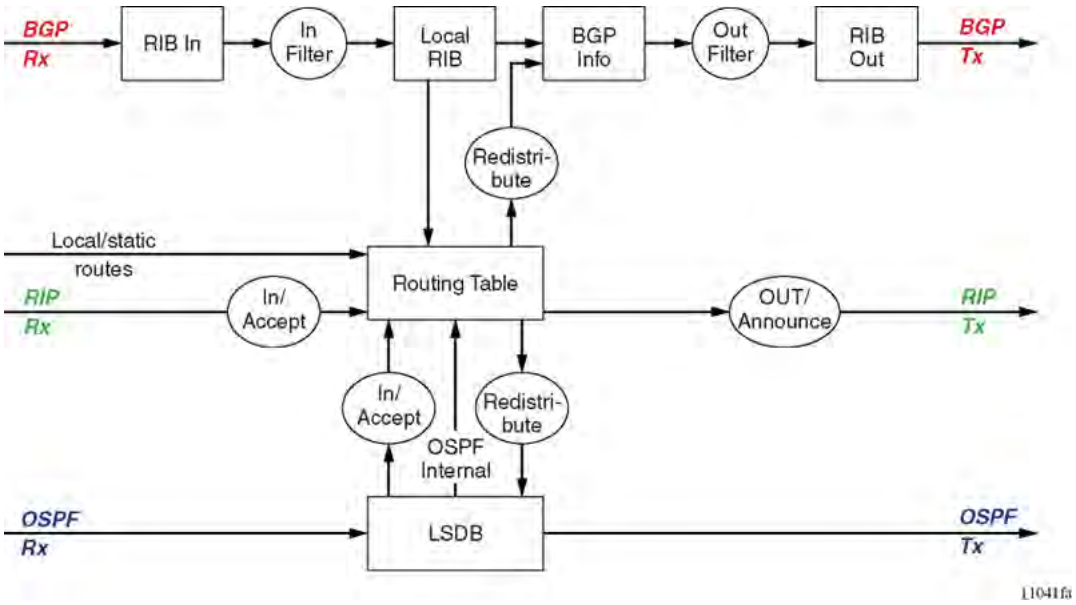


Figure 5: Route filtering for BGP, RIP, and OSPF routing protocols

The following figure shows how filters apply to the IS-IS protocol for Avaya VENA Fabric Connect Layer 3 VSNs or IP Shortcuts.

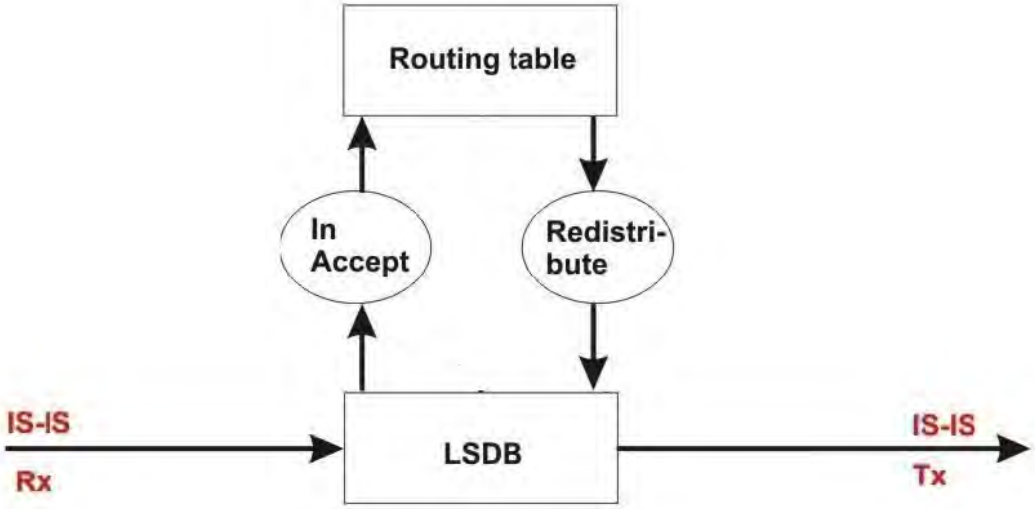


Figure 6: Route filtering for the IS-IS routing protocol

Accept policies

Accept policies are applied to incoming traffic to determine whether to add the route to the routing table. Accept policies are applied differently to protocols, as follows:

- RIP and BGP—filters apply to all incoming route information.

- OSPF—filters apply only to external route information. Internal routing information is not filtered because otherwise, other routers in the OSPF domain can have inconsistent databases that can affect the router view of the network topology.
- IS–IS —filters apply to all incoming route information.

In a network with multiple routing protocols, you can prefer specific routes from RIP instead of from OSPF. The network prefix is a commonly used match criterion for accept policies.

Redistribution filters

Redistribution filters notify changes in the route table to the routing protocol (within the device). With redistribution filters, providing you do not breach the protocol rules, you can choose not to advertise everything that is in the protocol database, or you can summarize or suppress route information. By default, no external routes are leaked to protocols that are not configured.

Announce policies

Announce policies are applied to outgoing advertisements to neighbors or peers in the protocol domain to determine whether to announce specific route information. Out filtering applies to RIP updates and BGP NLRI updates.

In contrast, announce policies are not applied to IS-IS or OSPF information because routing information must always be consistent across the domain. To restrict the flow of external route information in the IS-IS or OSPF protocol database, apply redistribution filters instead of announce policies.

Route filtering stages

The following figure shows the three distinct filter stages that are applied to IP traffic.

These stages are:

- Filter stage 1 is the accept policy or in filter that applies to incoming traffic to detect changes in the dynamic (protocol-learned) routing information, which are then submitted to the routing table.
- Filter stage 2 is the redistribution filter that applies to the entries in the routing table to the protocol during the leaking process.
- Filter stage 3 is the announce policy or out filter that applies to outgoing traffic within a protocol domain.

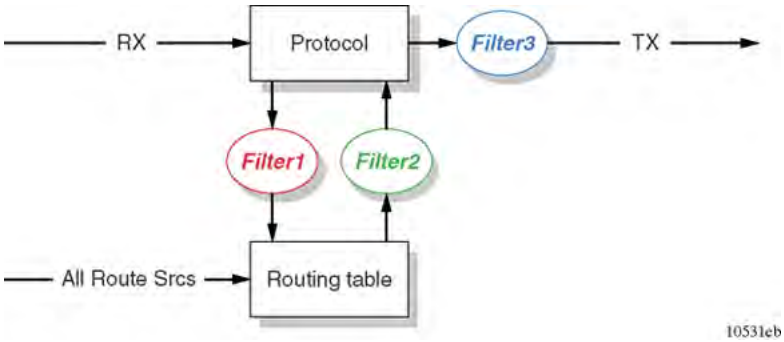


Figure 7: Route filtering stages

The following figure shows the logical process for route filtering on the switch.

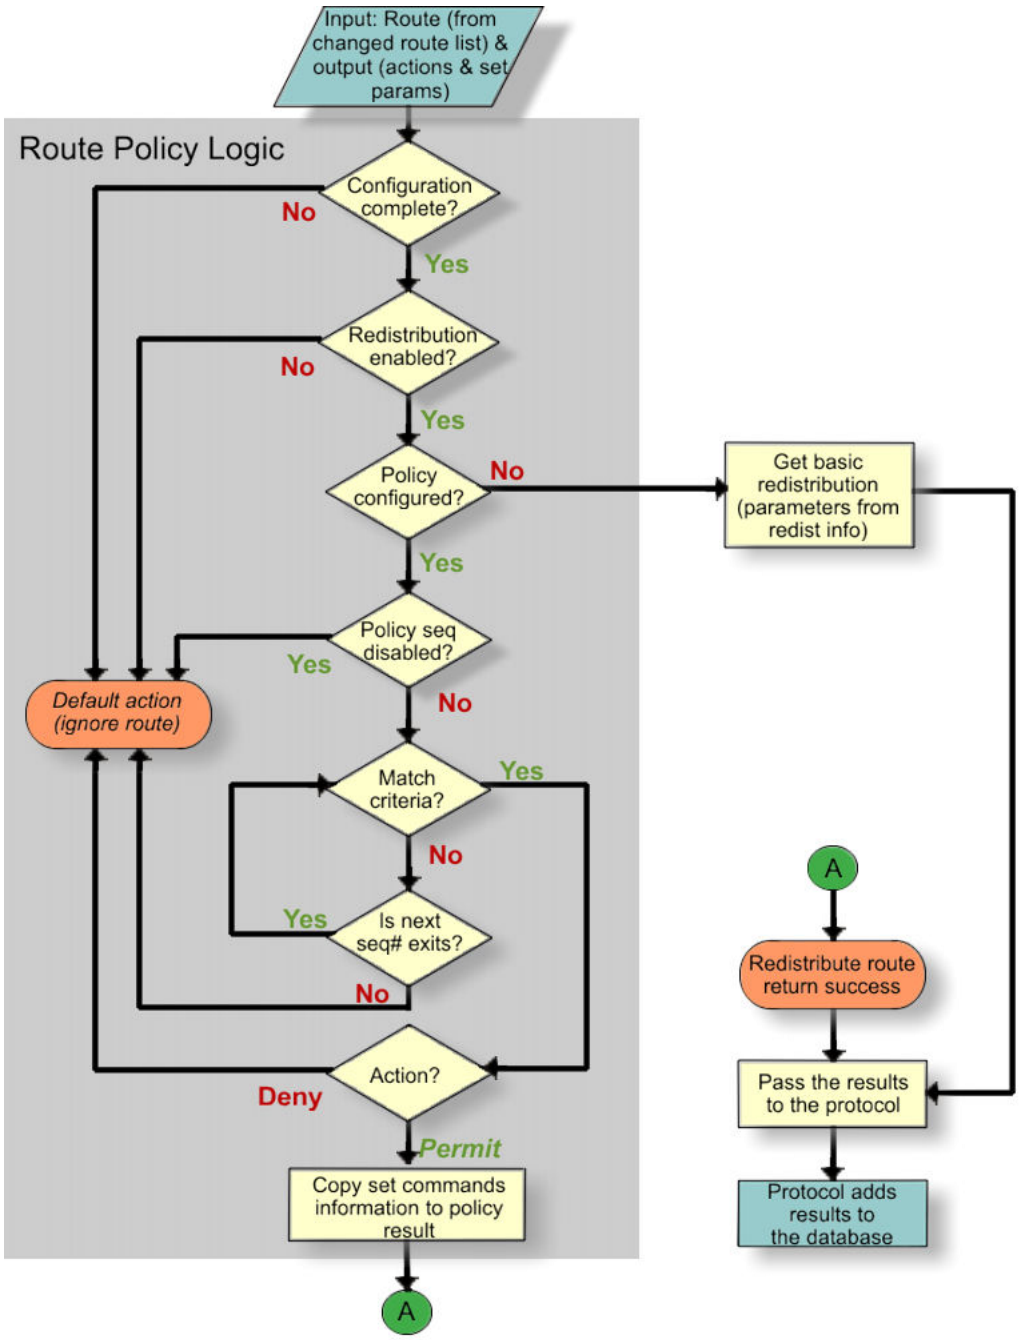


Figure 8: Route filtering logic

Prefix list

In the switch software, you can create one or more IP prefix lists and apply these lists to IP route policy.

Route policy definition

You can define an IP route policy and its attributes globally, and then apply them individually to interfaces and protocols. You can also form a unified database of route policies that the RIP or OSPF protocol can use for type of filtering purpose. A name or ID identifies a policy.

Under a policy you can have several sequence numbers. If you do not configure a field in a policy, the field appears as 0 in ACLI show command output. This value indicates that the device ignores the field in the match criteria. Use the clear option to remove existing configurations for the field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference field, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce or redistribute purposes.

You can only apply one policy for each purpose (RIP Announce, for example) on a given RIP interface. In this case, all sequence numbers under the policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

The following tables display the accept, announce, and redistribute policies for RIP, OSPF, IS-IS and BGP. The tables also display which matching criteria apply for a certain routing policy. In these tables, 1 denotes advertise router, 2 denotes RIP gateway, and 3 denotes that external type 1 and external type 2 are the only options.

Table 3: Protocol route policy table for RIP

	Announce				Accept
	OSPF	Direct	RIP	BGP	RIP
Match Protocol	Yes	Yes	Yes	Yes	
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source	Yes[1]		Yes[2]		
Match NextHop	Yes	Yes	Yes	Yes	Yes
Match Interface			Yes		
Match Route Type	Yes				
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
NssaPbit					
SetRoute Preference					Yes
SetMetric TypeInternal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type					

Comments? infodev@avaya.com

	Announce				Accept
	OSPF	Direct	RIP	BGP	RIP
SetNextHop					
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					Yes
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 4: Protocol route policy table for OSPF

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
Match Protocol					
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes[2]		
Match NextHop		Yes	Yes	Yes	
Match Interface			Yes		
Match Route Type					Yes[3]
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
NssaPbit					
SetRoute Preference					Yes
SetMetric TypeInternal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	
SetNextHop				Yes	
SetInject NetList	Yes	Yes	Yes	Yes	Yes

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
SetMask					
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 5: Protocol route policy table for IS-IS

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
Match Protocol					
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source					
Match NextHop		Yes	Yes	Yes	
Match Interface			Yes		
Match Route Type					Yes[3]
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
NssaPbit					
SetRoute Preference					Yes
SetMetric TypeInternal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	
SetNextHop				Yes	
SetInject NetList					
SetMask					
SetAsPath					

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 6: Protocol route policy table for BGP

	Redistribute			Accept	Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
Match as-path				Yes	Yes
Match community	Yes	Yes	Yes	Yes	Yes
Match community-exact				Yes	Yes
Match extcommunity				Yes	Yes
Match interface					
Match local-preference					
Match metric	Yes	Yes	Yes	Yes	Yes
Match network	Yes	Yes	Yes	Yes	Yes
Match next-hop		Yes	Yes	Yes	Yes
Match protocol					
Match route-source				Yes	
Match route-type			Yes		Yes
Match tag					
Match vrf					
Match vrfids					
Set as-path				Yes	Yes
Set as-path-mode				Yes	Yes
Set automatic-tag					
Set community				Yes	Yes
Set community-mode				Yes	Yes
Set injectlist	Yes	Yes	Yes		
Set ip-preference					

	Redistribute			Accept	Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
Set local-preference				Yes	Yes
Set mask					
Set metric	Yes	Yes	Yes	Yes	Yes
Set metric-type					
Set metric-type-internal					
Set next-hop				Yes	Yes
Set nssa-pbit					
Set origin					Yes
Set origin-egp-as					
Set Tag					
Set Weight				Yes	

Individual port routing control

You can enable or disable routing capabilities on specified device ports, even when the port is part of a routed VLAN. For example, after you disable IP routing on a specific port, the IP traffic that enters that port is not routed to another interface on the device.

You can use this feature as a security measure to prevent untrusted VLAN ports from injecting IP traffic that is destined to be routed by the device. This feature only stops routed protocols, for example, RIP or OSPF, not static routing.

Reverse path checking

When enabled, reverse path checking (RPC) prevents packet forwarding for incoming IP packets that have incorrect or forged (spoofed) IP addresses. RPC guarantees that traffic received on one interface was sent by a station from the identified interface, which prevents address spoofing. The Avaya Virtual Services Platform 9000 performs a reverse path check to determine if the packet IP address is verifiable. If the IP address is not verifiable, the packet is discarded.

You configure RPC for each IP interface. When RPC is enabled, the Avaya Virtual Services Platform 9000 checks all routing packets that come through that interface. The system ensures that the source address and source interface appear in the routing table, and that the value matches the interface on which the packet was received.

You can use one of two modes for RPC:

- Exist-only mode: In this mode, RPC checks whether the source IP address for the incoming packet exists in the routing table. If the source IP entry is found, the packet is forwarded as usual; otherwise, the packet is discarded.
- Strict mode: In this mode, RPC checks whether the source IP address for the incoming packet exists in the routing table. If the source IP entry is found, RPC further checks if the source IP interface address matches the packet incoming interface address. If they match, the packet is forwarded as usual, otherwise, the packet is discarded.

The following example illustrates how strict mode for reverse path checking works.



Figure 9: Reverse path checking network configuration

Consider the following parameters:

- A router connects a server (32.57.5.10) to a client (192.32.45.10).
- The router uses reverse path checking.
- The router has the following entries in the routing table:

Table 7: Routing table

Destination address	Next-hop address	Forward through port
32.57.5.10	173.56.42.2	3/7
192.32.45.10	145.34.87.2	7/2
192.32.46.10	145.34.88.2	7/1

If the client sends a legitimate packet, the following actions occur:

- The client sends a packet to the server. The packet has a source IP address of 192.32.45.10 and a destination IP address of 32.57.5.10.
- The packet arrives at router port 7/2 (brouter). The routing engine performs a destination IP address lookup and finds the destination port is 3/7.
- Reverse path checking begins. The routing engine searches for the source IP address of 192.32.45.10. The routing engine finds an entry in the routing table that specifies the next-hop port as 7/2, which matches the packet incoming port. Because the address and port information matches, the switch forwards the packet as usual.

If the client sends a spoofed packet, the following actions occur:

- The client sends a packet to the server with a forged IP address of 192.32.46.10 through port 7/2.
- Reverse path checking finds that the source IP address next-hop port is 7/1, which does not match the packet incoming port of 7/2. In this case, the switch discards the packet.

You can think about reverse path checking as follows. If A sends packets to B through route X ingress port Y, then the return packets from B to A must egress X through the same port Y. If returning packets take a different path, the switch drops them.

Address Resolution Protocol

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host IP address, the network station uses Address Resolution Protocol (ARP) to determine the physical address for a network host by binding a 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.

The network station uses ARP to determine the host physical address as follows:

- 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.
- All network hosts receive the broadcast request.
- Only the specified host responds with its hardware address.
- The network station then maps the host IP address to its physical address and saves the results in an address-resolution cache for future use.
- The network station ARP table displays the associations of the known MAC address to IP address.

You can create ARP entries, and you can delete individual ARP entries.

Enable ARP traffic

The Avaya Virtual Services Platform 9000 accepts and processes ARP traffic, spanning tree bridge packet data units (BPDU), and Topology Discovery Protocol packets on port-based VLANs with the default port action of drop. If a filter port action is drop for a packet, ARP packets are also dropped. As a result, ARP entries on that port are cleared and are not relearned when the ARP aging timer expires. To prevent dropped ARP packets, configure the following options:

- A user-defined protocol-based VLAN for ARP EtherType (byprotocol usrDefined 0x0806).
- Ports as static members to this VLAN with the default port action of drop.
- The port default VLAN ID to the correct port-based VLAN where the ARPs are processed.

You do not need to make configuration changes for the BPDU and Topology Discovery Protocol packets.

Only one user-defined protocol-based VLAN for ARP is allowed for each Spanning Tree Group (STG). If the ports with the default port action of drop are in different STGs, you must create additional user-defined protocol-based VLANs.

Proxy ARP

A network station uses proxy ARP to respond to an ARP request from a locally attached host or end station for a remote destination. The network station sends an ARP response back to the local host with its own MAC address of the network station interface for the subnet on which the ARP request was received. The reply is generated only if the device has an active route to the destination network.

The following figure shows an example of proxy ARP operation. In this example, host C with mask 24 appears to be locally attached to host B with mask 16, so host B sends an ARP request for host C. However, the Avaya Virtual Services Platform 9000 is between the two hosts. To enable communication between the two hosts, the Avaya Virtual Services Platform 9000 responds to the ARP request with the IP address of host C but with its own MAC address.

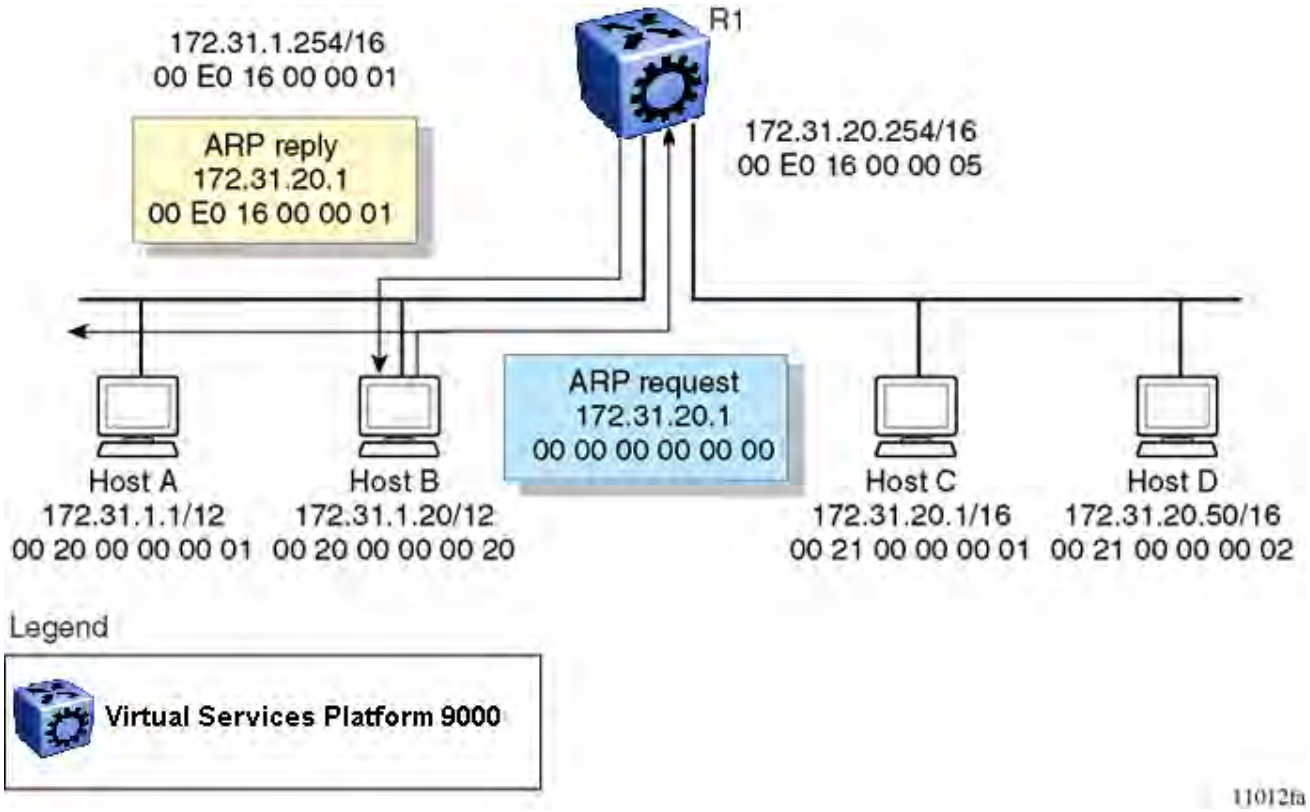


Figure 10: Proxy ARP operation

Loop detection

To prevent cases of ARP looping, configure the ARP loop detection flag to detect this situation. When a loop is detected, the port is shut down.

Flushing router tables

For administrative or troubleshooting purposes, sometimes you must flush the routing tables. Flush routing tables either by VLAN or by port. In a VLAN context, all entries associated with the VLAN are flushed. In a port context, all entries associated with the port are flushed.

Reverse Address Resolution Protocol

Certain devices use the Reverse Address Resolution Protocol (RARP) to obtain an IP address from a RARP server. MAC address information for the port is broadcast on all ports associated with an IP protocol-based or port-based VLAN. To enable a device to request an IP address from a RARP server outside its IP VLAN, you must create a RARP protocol-based VLAN.

RARP has the format of an ARP frame but its own Ethernet type (8035). You can remove RARP from the IP protocol-based VLAN definition and treat it as a separate protocol, thus creating a RARP protocol-based VLAN.

A typical network topology provides desktop switches in wiring closets with one or more trunk ports that extend to one or more data center switches where attached servers provide file, print, and other services. Use RARP functionality to define all ports in a network that require access to a RARP server as potential members of a RARP protocol-based VLAN. You must define all tagged ports and data center RARP servers as static or permanent members of the RARP VLAN. Therefore, a desktop host broadcasts an RARP request to all other members of the RARP VLAN. In normal operation, these members include only the requesting port, tagged ports, and data center RARP server ports. Because all other ports are potential members of this VLAN and RARP is only transmitted at startup, all other port VLAN memberships expire. With this feature, one or more centrally located RARP servers extend RARP services across traditional VLAN boundaries to reach desktops globally.

DHCP option 82

The DHCP option 82 is the DHCP Relay Agent Information option. The DHCP relay agent inserts option 82 when it forwards the client-originated DHCP packets to a DHCP server. The Relay Agent Information option is organized as a single DHCP option that contains one or more sub-options that convey information known by the relay agent. The DHCP server echoes the option back to the relay agent in server-to-client replies, and the relay agent removes the option before forwarding the reply to the client.

The DHCP option 82 is added at the DHCP relay level as shown in the following image.

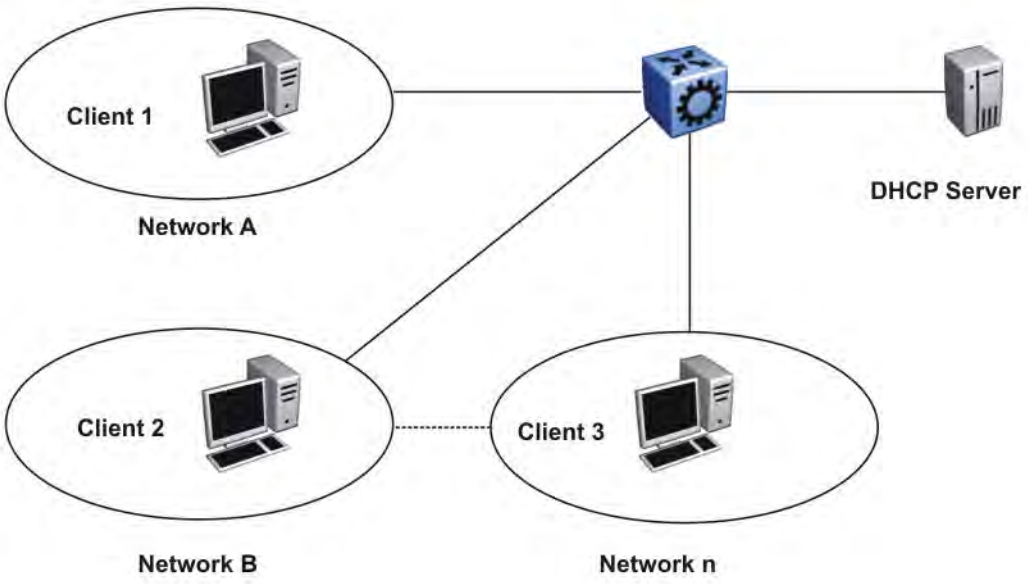


Figure 11: DHCP Client-Relay-Server Architecture

The Relay Agent Information option (code 82) is a container for specific agent-supplied suboptions; Agent Circuit ID (code 1) and Agent Remote ID (code 2). The suboptions can represent different information relevant for the relay. The fields are encoded in the following manner, where N or n is the total number of octets in the Agent Information Field (all bytes of the suboptions):

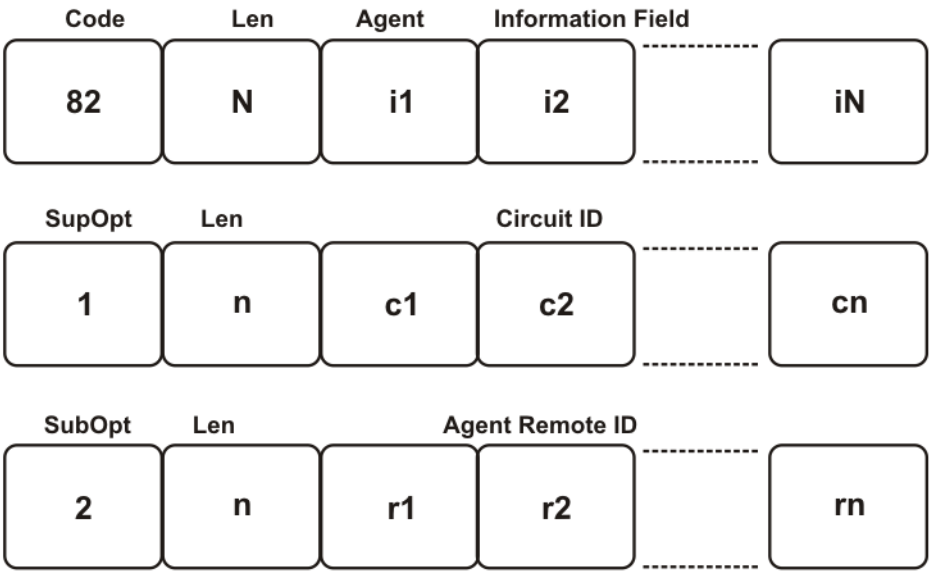


Figure 12: Format of the Relay Agent Information

Because at least one of the sub-options must be defined, the minimum Relay Agent Information length is two (2), and the length n of the suboption can be zero (0). The sub-options do not have to appear in any particular order. No pad suboption is defined and the Information field is not terminated with 255 suboption.

Suboptions

The suboptions are Agent Circuit ID and Agent Remote ID.

The DHCP relay agents can add the Agent Circuit ID to terminate switched or permanent circuits. The Agent Circuit ID encodes an agent-local identifier of the circuit from which a DHCP client-to-server packet was received. Agents can use the Circuit ID to relay DHCP responses back to the proper circuit. In the Avaya Virtual Services Platform 9000, the Agent Circuit ID field contains the ifIndex of the interface on which the packet is received.

DHCP relay agents can add the Agent Remote ID to terminate switched or permanent circuits, and can identify the remote host end of the circuit. The Avaya Virtual Services Platform 9000 uses the Agent Remote ID field to encode the MAC address of the interface on which the packet is received. The Agent Remote ID must be globally unique.

Agent operations

A DHCP relay agent adds a Relay Agent Information field as the last option in the DHCP options field of any recognized BOOTP or DHCP packet forwarded from a client to a server. However, if the End Option 255 is present, then the DHCP relay agent adds a Relay Agent information field before the End Option 255 field.

Relay agents can receive a DHCP packet from an untrusted circuit with the gateway IP address (GIADDR) set to zero to indicate that the relay agent is the first-hop router from the gateway. If a Relay Agent Information option is present in the packet, the relay agent discards the packet and increments an error counter. A trusted circuit can contain a trusted downstream network element, for example, a bridge, between the relay agent and the client. The bridge can add a relay agent option but does not set the GIADDR field. In this case, the relay agent forwards the DHCP packet per normal DHCP relay agent operations, and sets the GIADDR field to the relay address. The relay agent does not add a second relay agent option.

You can distinguish between a trusted circuit and an untrusted circuit based on the type of circuit termination equipment you use. To make a circuit trusted, set the trusted flag under DHCP for each interface.

After packets append the Relay Agent Information option, the packets that exceed the MTU or the vendor size buffer of 64 bits, are forwarded without adding the Agent Information option, and an error counter is incremented.

The relay agent or the trusted downstream network element removes the Relay Agent Information option echoed by a server that is added when forwarding a server-to-client response back to the client.

The following list outlines the operations that the relay agent does not perform:

- The relay agent does not add an Option Overload option to the packet or use the file or sname fields to add the Relay Agent Information option. The agent does not parse or remove Relay Agent Information options that can appear in the sname or file fields of a server-to-client packet forwarded through the agent.
- The relay agent does not monitor or modify client-originated DHCP packets addressed to a server unicast address; this includes the DHCP-REQUEST sent when entering the RENEWING state.
- The relay agent does not modify DHCP packets that use the IPSEC Authentication Header or IPSEC Encapsulating Security Payload.

A DHCP relay agent can receive a client DHCP packet forwarded from a BOOTP/DHCP relay agent closer to the client. This packet has a GIADDR as non-zero, and may or may not already have a DHCP Relay Agent option in it.

Relay agents configured to add a Relay Agent option which receive a client DHCP packet with a nonzero GIADDR, discards the packet if the GIADDR spoofs a GIADDR address implemented by the local agent itself. Otherwise, the relay agent forwards any received DHCP packet with a valid non-zero GIADDR without adding any relay agent options. The GIADDR value does not change.

UDP broadcast forwarding

Some network applications, such as the NetBIOS name service, rely on a User Datagram Protocol (UDP) broadcast to request a service or locate a server for an application. If a host is on a network, subnet segment, or VLAN that does not include a server for the service, UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. You can resolve this problem by forwarding the broadcasts to the server through physical or virtual router interfaces.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface out to other router IP interfaces as a rebroadcast or to a configured IP address. If the address is that of a server, the packet is sent as a unicast packet to this address. If the address is that of an interface on the router, the frame is rebroadcast.

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

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

For each ingress interface and protocol, the policy specifies how the UDP broadcast is retransmitted: to a unicast host address or to a broadcast address.

Virtual Router Redundancy Protocol

Because end stations often use 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 when the single static default gateway router for an end station is lost. VRRP introduces a virtual IP address (transparent to users) shared between two or more routers that connect the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides dynamic default gateway redundancy in the event of failover.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

The VRRP router that controls the IP addresses associated with a virtual router is the primary router and it forwards packets to these IP addresses. The election process provides a dynamic transition of forwarding responsibility if the primary router becomes unavailable.

In the following figure, the first three hosts install a default route to the R1 (virtual router 1) IP address and the other three hosts install a default route to the R2 (virtual router 2) IP address.

This configuration not only shares the load of the outgoing traffic, but it also provides full redundancy. If either router fails, the other router assumes responsibility for both addresses.

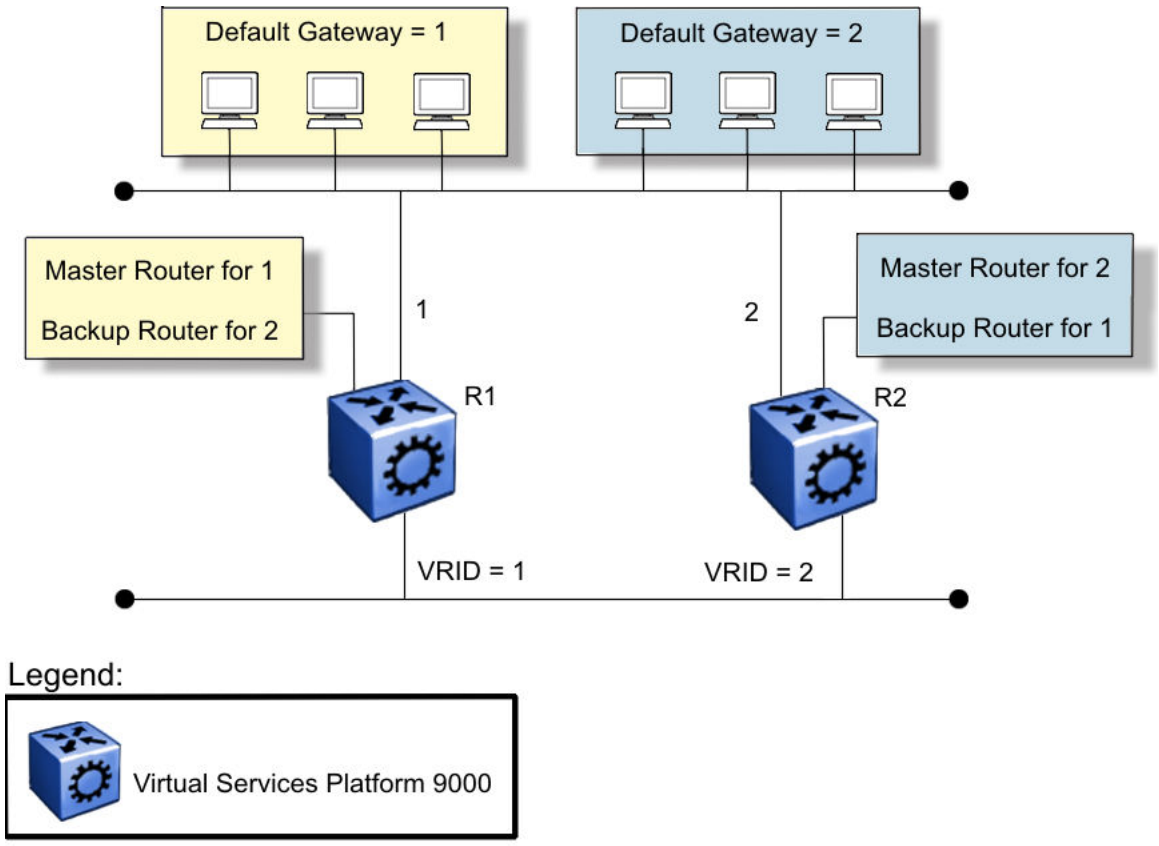


Figure 13: Virtual Router Redundancy Protocol configuration

The Avaya Virtual Services Platform 9000 supports 256 VRRP interfaces for each VRF and 512 VRRP interfaces for each system. The following terms are specific to VRRP:

- VRRP router—a router running the VRRP protocol
- Virtual router—an abstract object acting as the default router for one or more hosts, consisting of a virtual router ID and a set of addresses
- IP address owner—the VRRP router that has virtual router IP addresses as real interface addresses (This router responds to packets sent to this IP address.)
- Primary IP address—an IP address selected from the real addresses and used as the source address of packets sent from the router interface (The virtual primary router sends VRRP advertisements using this IP address as the source.)
- Virtual primary router—the router that assumes responsibility to forward packets sent to the IP address associated with the virtual router and answer ARP requests for these IP addresses
- Virtual router backup—the virtual router that becomes the primary router if the current primary router fails

When a VRRP router is initialized, if it is the IP address owner, its priority is 255 and it sends a VRRP advertisement. The VRRP router also broadcasts a gratuitous ARP request that contains the virtual router MAC address for each IP address associated with the virtual router. The VRRP router then transitions to the controlling state.

In the controlling state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. The VRRP router responds to ARP requests for these IP

addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts only packets addressed to IP addresses associated with the virtual router if it is the IP address owner. If the priority is not 255, the router transitions to the backup state to ensure that all Layer 2 switches in the downstream path relearn the new origin of the VRRP MAC addresses.

In the backup state, a VRRP router monitors the availability and state of the primary 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 backup router transitions back to the initialize state. If the primary router goes down, the backup router sends the VRRP advertisement and ARP request described in the preceding paragraph and transitions to the controlling state.

Whenever a packet is redirected on the same IP subnet on which it is received, Virtual Services Platform 9000 sends an Internet Control Message Protocol (ICMP) redirect packet data unit (PDU) to the IP address source of the packet. ICMP redirect uses the VRRP IP subnet as the source IP address for the end stations using the VRRP IP address as the next hop.

If an advertisement timer becomes active, the router sends an advertisement. If an advertisement is received with a 0 priority, the router sends an advertisement. The router transitions to the backup state in the following situations:

- If the priority is greater than the local priority
- If the priority is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address

Otherwise, the router discards the advertisement. If a shutdown occurs, the primary router sends a VRRP advertisement with a priority of 0 and transitions to the initialize state.

Critical IP address

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

As a result, if the local router IP interface connecting the virtual router to the external network fails, this does not automatically trigger a master router failover.

The critical IP address resolves this issue. If the critical IP address fails, it triggers a failover of the master router.

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

In VRRP, the local network uplink interface on router 1 is shown as the critical IP address for router 1. As well, the same network uplink is shown as the critical IP address for router 2. Router 2 also requires a critical IP address for cases in which it assumes the role of the master router.

With the support of VRRP and the critical IP interface linked to VRRP, you can build reliable small core networks that provide support for converged applications, such as voice and multimedia.

VRRP and SMLT

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

A deficiency occurs when VRRP-enabled switches use Split MultiLink Trunking (SMLT). If VRRP switches are aggregated into two Split MultiLink Trunk switches, the end host traffic is load-shared on all uplinks to the aggregation switches (based on the Multilink Trunk traffic distribution algorithm).

However, VRRP usually has only one active routing interface enabled. All other VRRP routers are in backup mode. Therefore, all traffic that reaches the backup VRRP router is forwarded over the interswitch trunk (IST) link toward the master VRRP router. In this case, the IST link does not have enough bandwidth to carry all the aggregated traffic.

To resolve this issue, assign the backup router as the backup master router. The backup master router can actively load-share the routing traffic with a master router.

When the backup master router is enabled, the incoming host traffic is forwarded over the SMLT links as usual. When the backup master router is configured along with the critical IP interface and the critical IP interface goes down, the VRRP router transitions to be the backup router with the backup master state down. In this state, the VRRP router does not forward traffic.

The following figure shows a sample VRRP configuration with SMLT. Because Router B is the backup master, routing traffic is load-shared between the two devices.

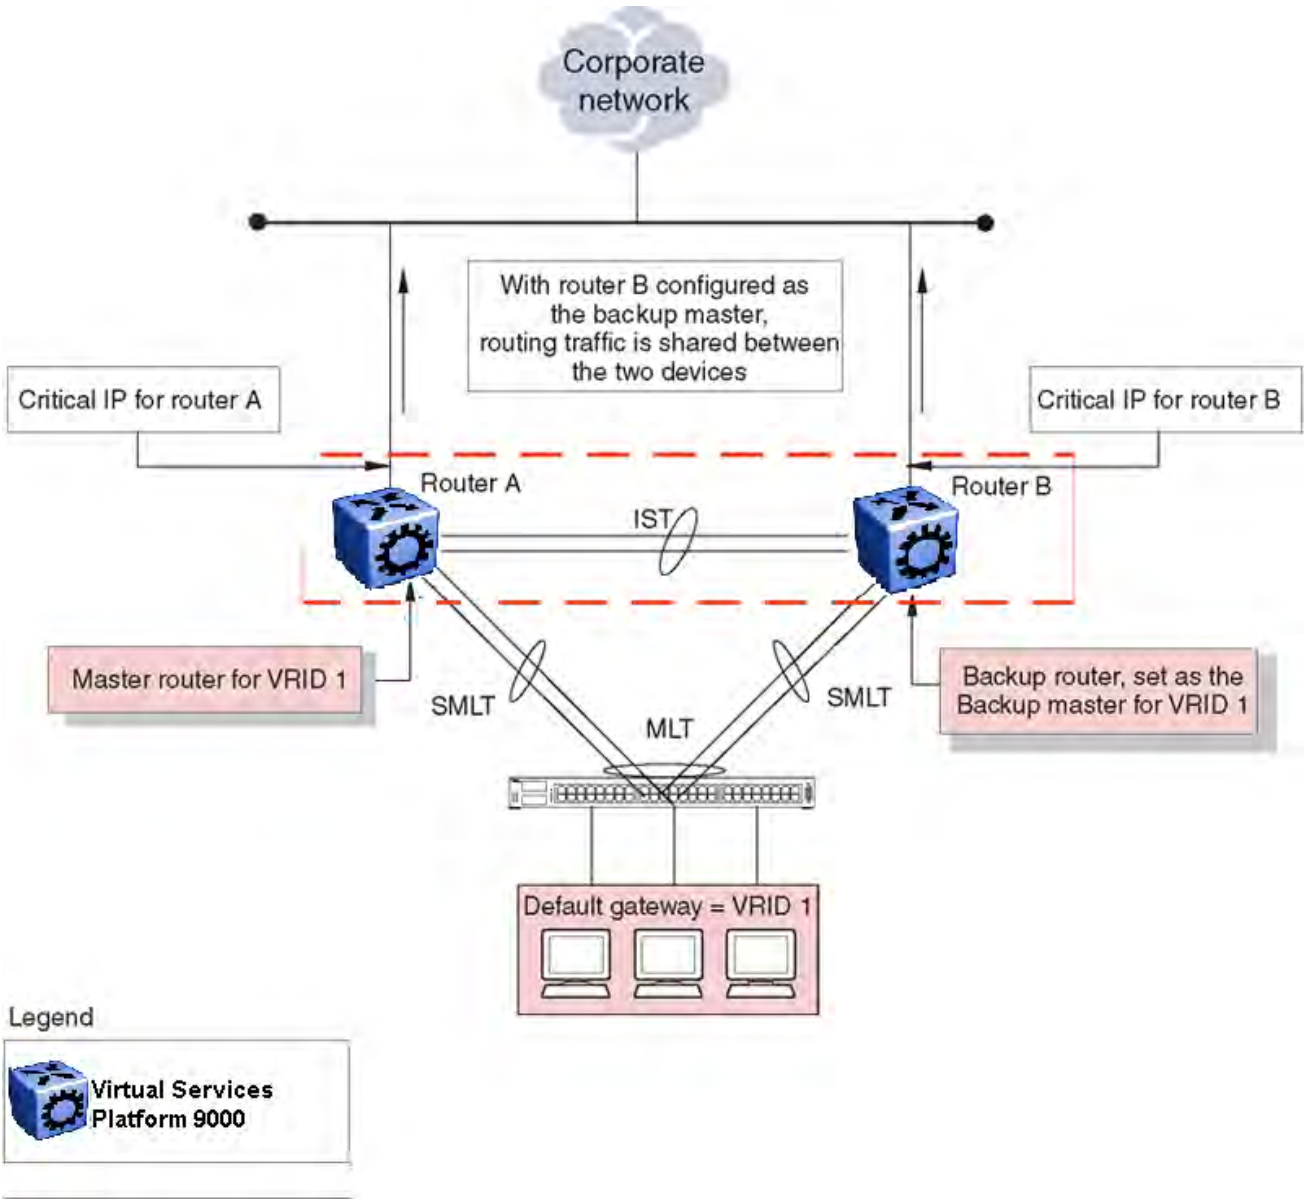


Figure 14: VRRP configuration with SMLT

VRRP fast hello timers

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

To meet these requirements, Avaya has two enhancements: Fast Advertisement Enable and Fast Advertisement Interval.

Fast Advertisement Enable acts like a toggle device for the Advertisement Interval and the Fast Advertisement Interval. When Fast Advertisement Enable is enabled, the Fast Advertisement Interval is used instead of the Advertisement Interval.

The Fast Advertisement Interval is similar to the current Advertisement Interval parameter except for the unit of measure and the range. The Fast Advertisement Interval is expressed in milliseconds and the range is from 200 to 1000 milliseconds. This unit of measure must be in multiples of 200 milliseconds, otherwise an error appears.

When you enable the fast advertisement interval, VRRP can only communicate with other Avaya Virtual Services Platform 9000 modules with the same settings.

RSMLT

In many cases, core network convergence time depends on the length of time a routing protocol requires to successfully converge. Depending on the specific routing protocol, this convergence time can cause network interruptions that range from seconds to minutes.

Avaya Routed Split MultiLink Trunking (RSMLT) permits rapid failover for core topologies by providing an active-active router concept to core SMLT networks.

RSMLT scenarios include SMLT triangles, squares, and SMLT full-mesh topologies, with routing enabled on the core VLANs.

Routing protocols include the following protocol types:

- IP Unicast Static Routes
- RIP1
- RIP2
- OSPF
- BGP

In the event of core router failures, RSMLT manages packet forwarding, thus eliminating dropped packets during the routing protocol convergence.

SMLT/RSMLT operation in Layer 3 environments

Figure 15: SMLT and RSMLT in Layer 3 environments on page 47 shows a typical redundant network example with user aggregation, core, and server access layers. To minimize the creation of many IP subnets, one VLAN (VLAN 1, IP subnet A) spans all wiring closets.

SMLT provides the loop-free topology and forwards all links for VLAN 1, IP subnet A.

The aggregation layer switches are configured with routing enabled and provide active-active default gateway functionality through RSMLT.

After you enable RSMLT on a VLAN (on both aggregation devices), the cluster devices simply inform each other (over IST messaging) of their physical IP or MAC on that VLAN. Thereafter, the two cluster devices take mutual ownership of their IP addresses on that VLAN. This action means each cluster device

- Replies to ARP requests for both the IP and the peer IP on that VLAN

- Replies to pings to the IP and the peer IP on that VLAN
- Routes IP traffic that is directed to the physical MAC of the IP or the physical MAC of the peer IP on that VLAN

In this case, routers R1 and R2 forward traffic for IP subnet A. RSMLT provides both router failover and link failover. For example, if the Split MultiLink Trunk link between R2 and R4 is broken, the traffic fails over to R1 as well.

For IP subnet A, VRRP with a backup master can provide the same functionality as RSMLT, as long as no additional router is connected to IP subnet A.

RSMLT provides superior router redundancy in core networks (IP subnet B), where OSPF is used for the routing protocol. Routers R1 and R2 provide router backup for each other, not only for the edge IP subnet A, but also for the core IP subnet B. Similarly routers R3 and R4 provide router redundancy for IP subnet C and also for core IP subnet B.

Router R1 failure

The following figure shows SMLT and RSMLT in Layer 3 environments.

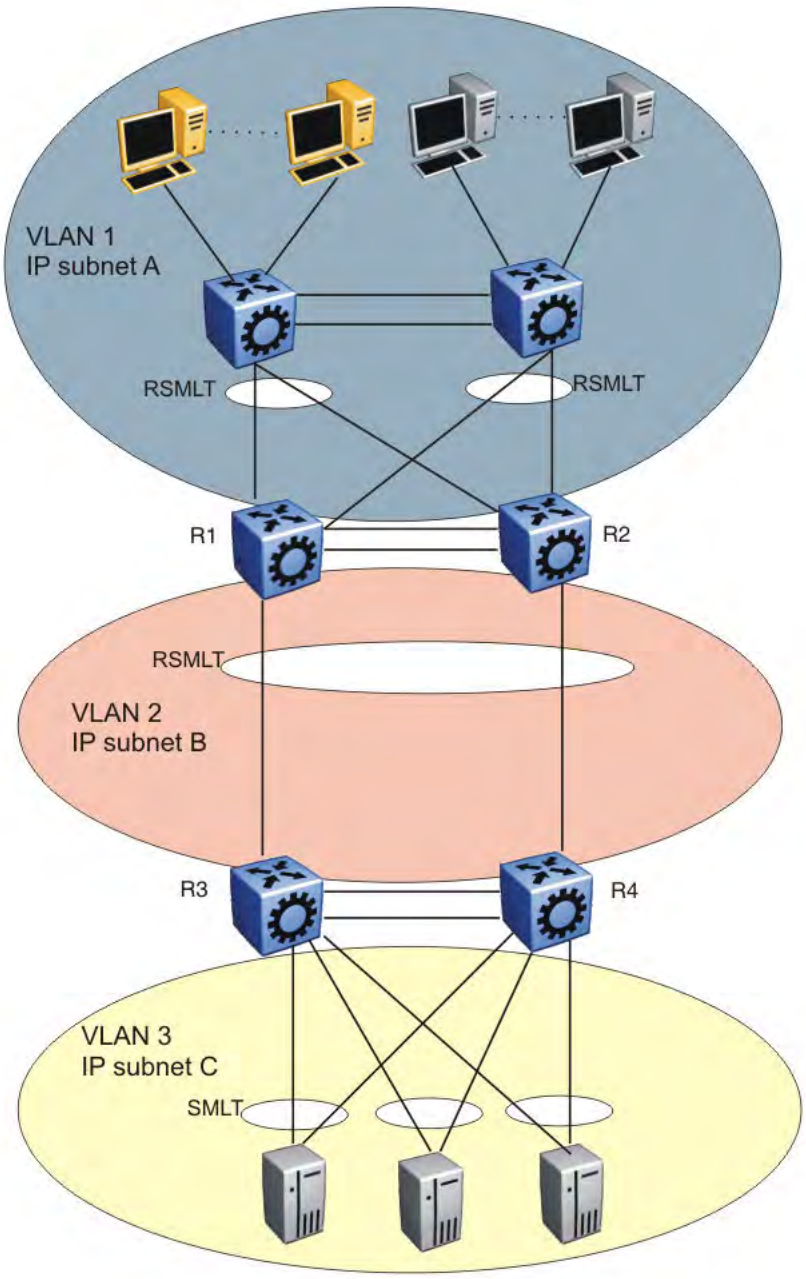


Figure 15: SMLT and RSMLT in Layer 3 environments

R3 and R4 both use R1 as their next hop to reach IP subnet A. Even though R4 sends the packets to R2, they are routed directly at R2 into subnet A. R3 sends its packets to R1 and they are also sent directly into subnet A. After R1 fails, all packets are directed to R2, with SMLT. R2 still routes for R2 and R1. After OSPF convergence, the routing tables in R3 and R4 change their next hop to R2 to reach IP subnet A. You can configure the hold-up timer (that is, for the amount of time R2 routes for R1 in a failure) for a time period greater than the routing protocol convergence, you can configure it as indefinite (that is, the members of the pair always route for each other).

Avaya recommends that you use an indefinite hold-up timer value for applications that use RSMLT at the edge instead of VRRP.

Router R1 recovery

After R1 restarts after a failure, it becomes active as a VLAN bridge first. Packets destined to R1 are switched, using the bridging forwarding table, to R2 for as long as the hold-down timer is configured. Those packets are routed at R2 for R1. Similar to VRRP, the hold-down timer value must be greater than the time the routing protocol requires to converge its tables.

After the hold-down time expires and the routing tables converge, R1 starts routing packets for itself and also for R2. Therefore, it does not matter which of the two routers is used as the next hop from R3 and R4 to reach IP subnet A.

If single-homed IP subnets are configured on R1 or R2, Avaya recommends that you add another routed VLAN to the interswitch trunks (IST) with lower routing protocol metrics as a traversal VLAN/subnet to avoid unnecessary ICMP redirect generation messages. This recommendation also applies to VRRP implementations.

RSMLT network design and configuration

Because RSMLT is based on SMLT, all SMLT configuration rules apply. In addition, RSMLT is enabled on the SMLT aggregation switches for each VLAN. The VLAN must be a member of SMLT links and the IST trunk. For more information about how to configure SMLT in a Layer 2 environment, see *Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000,* NN46250-503.

The VLAN also must be routable (IP address configured) and you must configure an Interior Gateway Protocol (IGP) such as OSPF on all four routers, although it is independent of RSMLT. All routing protocols, even static routes, work with RSMLT.

The RSMLT pair switches provide backup for each other. As long as one of the two routers of an IST pair is active, traffic forwarding is available for both next hops R1/R2 and R3/R4.

RSMLT edge support

VSP 9000 stores the peer MAC and IP address pair in its local configuration file and restores the configuration if the peer does not restore after a simultaneous restart of both RSMLT-peer switches.

The RSMLT edge support feature adds an enhancement whereby the peer MAC (for the IP on the VLAN) is committed to the config.cfg file after you use the **`save config`** command. If you power off both devices, and then power up only one of them, that single device can still take ownership of its peer IP on that VLAN even if it has not seen that peer switch since it started. This enhancement is necessary if you configure the peer (the device which is still down) IP as the default gateway in end stations.

If you enable RSMLT edge support, you must also ensure that the hold-up timer for RSMLT on those edge VLANs equals infinity (9999). This timer value ensures that if one cluster device fails, the remaining cluster device maintains ownership of the failed peer IP indefinitely.

The edge VLAN can be tagged over SMLT links, single attached links, or more SMLT links.

Important:

If you clear the peer information the device can stop forwarding for the peer.

RSMLT implementation does not use a virtual IP address but instead uses physical IP addresses for redundancy. At the same time, you can deploy RSMLT in either routed configurations, or edge

configurations, where you previously used VRRP (and back-up master). Previously, if a power outage occurred or a shutdown of both switches within a dual core IST pair, only one device came back up. Clients using the powered-off device IP/MAC as the default gateway lost connectivity to the network. In such a scenario, even with RSMLT enabled on the device, it cannot act as a backup for the peer as it was unaware of the peer IP or MAC address.

After both the dual core IST switches come back, the IST is operational. If an RSMLT peer-enabled message is received from the peer, normal RSMLT operation occurs.

If the peer has either an IP or MAC change, you must save the configuration for the RSMLT edge support to operate correctly. However, if the IST peer up message is not received (for example, if you do not enable RSMLT properly), and you enable the RSMLT edge support flag, the RSMLT hold-down timer starts and permits routing protocols to converge; during this time user operation can be affected. After the hold-down timer expires, saved peer information is picked up and the device starts to act as backup for the peer by adding the previously saved MAC and ARP records.

The hold-up timer starts and after this timer expires the previously added MAC and ARP records are deleted and the device stops acting as backup for the peer, as the peer is not running proper RSMLT for the VLAN. The RSMLT is a parameter for each VLAN, and therefore all affects are on an individual VLAN basis, not necessarily a global device. Edge support mode uses the local values of the hold-down timer (default value of 60 seconds) and hold-up timer (default value of 180 seconds).

Chapter 4: ARP configuration using ACLI

Network stations that use IP protocol require both a physical address and an IP address to transmit packets. In situations where the station knows only the network host IP address, the Address Resolution Protocol (ARP) lets you use the network station to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address.

A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

ARP response is enabled by default.

Enabling ARP on a port or a VLAN

Before you begin

- You must log on to the VLAN, or GigabitEthernet Interface Configuration mode in ACLI.

About this task

Enable ARP on the device so that it answers local ARP requests.

You can enable or disable ARP responses on the device. You can also enable ARP proxy, which lets a router answer a local ARP request for a remote destination.

Procedure

Enable ARP on the device:

```
ip arp-response
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 200
VSP-9012:1(config-if)#ip arp-response
```

Enabling ARP proxy

Before you begin

- You must log on to Access the VLAN, or GigabitEthernet Interface Configuration mode in ACLI.

About this task

Configure an ARP proxy to allow the platform to answer a local ARP request for a remote destination. ARP proxy is disabled by default.

Procedure

Enable ARP proxy on the device:

```
ip arp-proxy enable
```

Use the `no` operator to disable ARP proxy: `no ip arp-proxy [enable]`

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 200
VSP-9012:1(config-if)#ip arp-proxy enable
```

Configuring ARP loop detection

Configure loop detect to determine if the same MAC address appears on different ports. Use the ARP-Detect feature to account for ARP packets on IP configured interfaces.

About this task

Loop detection works only after you enable loop-detect. To clear this option, you must disable the loop detection.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Configure ARP loop detection:

   ```
   loop-detect [action {port-down|mac-discard}] [arp-detect]
   ```

3. Configure ARP loop detection to the default:

   ```
   default loop-detect [action] [arp-detect]
   ```

4. Disable loop detection:

   ```
   no loop-detect [arp-detect]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 4/16
```

Configure loop detect to determine if the same MAC address appears on different ports and activate ARP-Detect to detect Layer 3 loops:

```
VSP-9012:1(config-if)#loop-detect action mac-discard arp-detect
```

Variable definitions

Use the data in the following table to use the **loop-detect** command.

Table 8: Variable definitions

Variable	Value
action	Indicates the action that the device takes: • port-down • mac-discard
arp-detect	The ARP-detect feature is used for IP configured interfaces for ARP packets. Enable this feature (in addition to loop detection) on route interfaces.

Showing ARP information

When you use the interface parameter with the **show ip arp** command you can display ARP configuration information only for a specific VSP 9000 switch.

The **show ip arp** command displays all of the configured and dynamically learned ARP entries in the ARP table.

Procedure

1. Display ARP information for a specified port or for all ports:

   ```
   show ip arp interface [gigabitethernet {slot/port[-slot/port][,...]]
   ```

2. Display ARP information for a VLAN:

   ```
   show ip arp interface [vlan <1-4084>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#show ip arp interface vlan 1

================================================================================
                                    Vlan Arp
================================================================================
VLAN ID  DOPROXY    DORESP
--------------------------------------------------------------------------------
1        false      true
```

Variable definitions

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

Table 9: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the IP address of a network.
gigabitethernet *{slot/port[-slot/port][,...]}*	Displays ARP entries for a particular brouter port.
interface	Displays ARP interface configuration information.
-s <*A.B.C.D*><*A.B.C.D*>	Specifies a subnet. You must indicate the IP address followed by the subnet mask expressed as <A.B.C.D> <A.B.C.D>.
spbm-tunnel-as-mac	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
static-mcastmac	Displays static multicast MAC ARP information.
vlan <*1–4084*>	Displays ARP entries for a particular VLAN ID, expressed as a value from 1 to 4084.
vrf *WORD*<*1–16*>	Specifies a VRF name expressed as text from 1 to 16 characters in length. The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis, including all VRFs (which includes the Management Router VRF).
vrfids *WORD*<*0–512*>	Specifies a range of VRFIDs as text from 0 to 512 characters in length. The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis, including all VRFs (which includes the Management Router VRF).

Use the data in the following table to help you understand the **show ip arp interface** command output.

Table 10: Variable definitions

Variable	Value
PORT_NUM	Indicates the port number.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.

Use the data in the following table to help you understand the **`show ip arp interface vlan`** command output.

Table 11: Variable definitions

Variable	Value
VLAN_ID	Indicates the VLAN ID.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.

Configuring IP ARP static entries

Configure ARP static entries to modify the ARP parameters on the device. The only way to change a static ARP is to delete the static ARP entry and create a new entry with new information.

Note:

Virtual Services Platform 9000 supports static ARP entries for NLB multicast and NLB multicast IGMP. Virtual Services Platform 9000 does not support static ARP entries for NLB unicast.

Before you begin

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

Procedure

Configure ARP static entries on the device:

```
ip arp <A.B.C.D> 0x00:0x00:0x00:0x00:0x00:0x00 {slot/port[-slot/port]
[,...]}
```

Example

Add ARP entries and configure static multicast MAC entries:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

```
VSP-9012:1(config)#ip arp 192.0.2.128 00-16-76-7D-80-C2 4/15
VSP-9012:1(config)#ip arp static-mcast 192.0.2.128 00-16-76-7D-80-C2 vid 200
```

Variable definitions

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

Table 12: Variable definitions

Variable	Value
multicast-mac-flooding [enable]	Determines whether ARP entries for multicast MAC addresses are associated with the VLAN or the port interface on which they were learned. Use the no operator to delete a static entry from the ARP table: `no ip arp multicast-mac-flooding [enable]` To configure this option to the default value, use the default operator with this command.
request-threshold <50-1000>	Configures the maximum number of outstanding ARP requests that a device can generate. The range is 50–1000. The default value is 500. To configure this option to the default value, use the default operator with this command.
static-mcast	Configures static multicast MAC entries.
timeout <1-32767>	Configures the length of time in seconds an entry remains in the ARP table before timeout. The range is 1–32767. To configure this option to the default value, use the default operator with this command.
<A.B.C.D>[0x00:0x00:0x00:0x00:0x00:0x00:0x00]{slot/port[-slot/port][,...]}	Adds ARP entries. *[0x00:0x00:0x00:0x00:0x00:0x00:0x00]* specifies the MAC address in hexadecimal format. The MAC address parameter does not accept MAC addresses beginning 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive). *{slot/port[-slot/port][,...]}* specifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports 93/2–3/4), or a series of slot and ports (3/2,5/36/2).

Clearing ARP entries

Before you begin

- You must log on to the Privileged EXEC mode in ACLI.

About this task

Use this procedure to clear dynamic ARP table entries associated with the interface or VLAN.

Procedure

Clear ARP entries:

```
clear ip arp interface <gigabitethernet|vlan> <slot/port[-slot/port]
[,...]|1-4084>
```

Example

```
VSP-9012:1>enable
```

Clear ARP entries:

```
VSP-9012:1#clear ip arp interface gigabitethernet 4/16
```

Variable definitions

Use the data in the following table to use the **clear ip arp interface** command.

Table 13: Variable definitions

Variable	Value
1–4084	Specifies the VLAN ID if you choose the VLAN interface type
gigabitethernet|vlan	Specifies the interface type
slot/port[-slot/port][,...]	Specifies the slot and port or range of slots and ports if you choose the fast Ethernet or Gigabit Ethernet interface type

Showing ARP table information

Show ARP information to view the configuration information in the ARP table.

About this task

When you use the interface parameter with the **show ip arp** command you can display ARP configuration information only for a specific VSP 9000 switch.

The**show ip arp** command displays all of the configured and dynamically learned ARP entries in the ARP table.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the ARP table:

```
show ip arp [<A.B.C.D>] [-s <A.B.C.D>] [gigabitEthernet <slot/port>]
[interface <gigabitethernet|vlan>] [spbm-tunnel-as-mac][static-
mcastmac <-s | vrf WORD<1-16>| vrfids WORD<0-512>| <A.B.C.D>][vlan
<1-4084>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-9012:show ip arp spbm-tunnel-as-mac

================================================================================
                              IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN    PORT  TYPE    TTL(10 Sec) TUNNEL

--------------------------------------------------------------------------------
10.20.10.20     00:00:00:00:00:01  0       -     LOCAL   2160
1.50.50.50      00:00:00:00:00:02  0       -     LOCAL   2160
10.255.255.255  ff:ff:ff:ff:ff:ff  0       -     LOCAL   2160
1.255.255.255   ff:ff:ff:ff:ff:ff  0       -     LOCAL   2160

================================================================================
                            IP Arp Extn - GlobalRouter
================================================================================
MULTICAST-MAC-FLOODING    AGING(Minutes)       ARP-THRESHOLD
--------------------------------------------------------------------------------
disable                   360                  500

4 out of 43 ARP entries displayed
```

Variable definitions

Use the data in the following table to help you use the **`show ip arp`** command.

Table 14: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the network IP address for the table.
-s <A.B.C.D>	Specifies the subnet for the table.
gigabitEthernet	Displays the entries for a particular brouter port.
interface	Displays ARP interface configuration information. Use the following paramaters to display ARP table information specifically for: • gigabitethernet {slot/port[–slot/port][,...]} displays IP ARP gigabitethernet interface information • VLAN <1–4084> displays IP ARP VLAN interface information Example: `show ip arp interface vlan 1`
spbm-tunnel-as-mac	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.

Variable	Value
static-mcasatmac	Displays static multicast MAC ARP information. Use these parameters to display ARP table information specifically for multicast MAC ARP as follows: • -s {A.B.C.D/X} —the specific ip/subnet value • vrf WORD<1–16>the static multicast MAC configurations for a particular VRF • vrfids WORD<0–512>—IP ARP static multicast MAC VRFIDs • {A.B.C.D}—specific network IP address
vlan	Displays ARP entries for a particular VLAN ID in a range from 1 to 4084. Use these parameters to display ARP table information specifically for: • vrf WORD<1–16>—the VLAN VRF name in a range from 1 to 16 characters • vrfids WORD<0–512>—the VLAN VRF ID in a range from 0 to 512 Example: `show ip arp vlan 1 vrf 1`
vrf *WORD <1-16>*	Specifies the name of the VRF. The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis including all VRFs (which includes the Mgmt Router VRF).
vrfids *WORD <0-512>*	Specifies the VRF ID. The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis including all VRFs (which includes the Mgmt Router VRF).

Use the data in the following table to help you understand the output of the **`show ip arp`** command.

Table 15: Variable definitions

Parameter	Description
IP_ADDRESS	Indicates the IP address where ARP is configured.
MAC_ADDRESS	Indicates the MAC address where ARP is configured.
VLAN	Indicates the VLAN address where ARP is configured.
PORT	Indicates the port where ARP is configured.
TYPE	Indicates the type of learning (dynamic or local) where ARP is configured.

Parameter	Description
TTL<10 secs>	Indicates the time to live as tenths of a second where ARP is configured.
TUNNEL	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
MULTICAST-MAC-FLOODING	Displays whether IP ARP multicast MAC flooding is enabled or disabled. When enabled, the ARP entries for multicast MAC addresses are associated with the VLAN or port interface on which they were learned.
AGING (Minutes)	Displays when the ARP aging timer expires.
ARP-THRESHOLD	Displays the maximum number of outstanding ARP requests that a device can generate.

Chapter 5: ARP configuration using Enterprise Device Manager

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host IP address, the network station can use Address Resolution Protocol (ARP) to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

Enabling or disabling ARP on the brouter port or a VRF instance

About this task

After you assign the IP address, you can configure ARP. By default, ARP Response is enabled and Proxy ARP is disabled.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **ARP** tab.
5. In the **DoProxy** check box, select **enable** to enable the Proxy ARP function.
6. In the **DoResp** check box, select **enable** to configure the system to respond to an ARP. The default is enable.
7. Click **Apply**.

   The ARP function is available only when the port or VLAN is routed; that is, it is assigned an IP address.

ARP field descriptions

Use the data in the following table to use the **ARP** tab fields.

Name	Description
DoProxy	Configures the system to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.

Enabling or disabling ARP on a VLAN or a VRF instance

About this task

To prevent dropped ARP packets, you must enable ARP on the VLAN before you enable ARP on the port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP**.
5. Click the **ARP** tab.
6. In the **DoProxy** field, click **enable** to enable the Proxy ARP function.
7. In the **DoResp** field, click **enable** to configure the system to respond to an ARP. The default is enable.
8. In the **DoFlood** field, click **enable** to configure the system to flood ARP responses for the Network Load Balancer Virtual MAC on the specified interface. The default is disable.
9. Click **Apply**.

   The ARP dialog box is available only if the port or VLAN is routed; that is, it is assigned an IP address.

ARP field descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
DoProxy	Configures the system to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.
DoFlood	Configures the system to flood ARP responses for the Network Load Balancer Virtual MAC on the specified interface. The default is disable.

Viewing and managing ARP

You can view and manage known MAC address to IP address associations. In addition, you can create or delete individual ARP entries.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **ARP** tab.
4. Click **Insert** to map an IP address to a physical address.
5. Type the IP address in the **NetAddress** box.
6. Select the port to which the address applies.
7. Type the physical media address in hexadecimal format in the **PhysAddress** box.
8. Click **Insert**.

ARP field descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
NetAddress	Specifies the IP address that corresponds to the physical address.
IfIndex	Specifies the interface to which this entry applies.
PhysAddress	Specifies the media-dependent physical address (that is, the Ethernet address).
Type	Specifies the type of ARP entry: • local—A locally configured ARP entry • static—A statically configured ARP entry • dynamic—A learned ARP entry • other—None of the preceding values

Name	Description
TimeToLive	Specifies the time-to-live, in seconds, for the ARP entry.
DestIfIndex	Configures the destination interface index to create a static ARP entry using SNMP.
DestVlanId	Specifies the destination VLAN ID.
Bmac	Specifies the Backbone MAC (B-MAC) address used to track where the NetAddress came from, if the entry is learned from an SPBM network.

Creating static ARP entries

Use the following procedure to create a static ARP entry.

Note:

Virtual Services Platform 9000 supports static ARP entries for NLB multicast and NLB multicast IGMP. Virtual Services Platform 9000 does not support static ARP entries for NLB unicast.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **ARP** tab.
4. Click **Insert**.
5. Click **Port**.

   OR

   Click **Port in VLAN**
6. In the dialog box, select the interface.
7. Click **OK**.
8. In the **IpAddress** field, type the IP address.
9. In the **MacAddress** field, type the MAC address.
10. Click **Insert**.

Configuring ARP proxy

About this task

With an ARP proxy, the Avaya Virtual Services Platform 9000 can respond to an ARP request from a locally attached host or end station for a remote destination. Proxy ARP does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet

on which the ARP request was received. The reply is generated only if the system has an active route to the destination network.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Choose a VLAN.
4. Click **IP**.
5. Click **ARP** tab.
6. Select **DoProxy enable**.
7. Click **Apply**.

Chapter 6: DHCP and UDP configuration using ACLI

Use Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), to provide host configuration information to the workstations dynamically. Use the DHCP relay commands to configure DHCP relay behavior on a port or on a VLAN.

This section describes ACLI commands for DHCP and User Datagram Protocol (UDP) configuration functions in Virtual Services Platform 9000.

Configuring DHCP parameters globally

Before you begin

- You must log on to the Global Configuration mode in ACLI.
- Configure an IP address on the interface to be used as the DHCP relay interface.

About this task

Configure DHCP relay parameters for the port or the VLAN.

Procedure

1. Create the forwarding path from the client to the server:

   ```
   ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D>
   ```

2. Enable the forwarding path from the client to the server:

   ```
   ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> enable
   ```

3. Modify DHCP mode to forward BootP messages only, DHCP messages only, or both:

   ```
   ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> mode <bootp|bootp_dhcp|
   dhcp>
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Create the forwarding path from the client to the server:

```
VSP-9012:1(config)#ip dhcp-relay fwd-path 43.17.159.120 43.17.121.50
```

Enable the forwarding path from the client the server:

```
VSP-9012:1(config)#ip dhcp-relay fwd-path 43.17.159.128 43.17.121.50
enable
```

Modify DHCP mode to forward both BootP and DHCP messagesy:

```
VSP-9012:1(config)#ip dhcp-relay fwd-path 43.17.159.128 43.17.121.50 mode
bootp_dhcp
```

Variable definitions

Use the data in the following table to use the **`ip dhcp-relay fwd-path`** command.

Table 16: Variable definitions

Variable	Value
fwd-path <A.B.C.D> <A.B.C.D>	Configures the forwarding path from the client to the server. A.B.C.D is the IP address configured on an interface (a locally configured IP address) to forward or relay BootP or DHCP. A.B.C.D is the IP address of the DHCP server in the network. If this IP address corresponds to the locally configured IP network, the DHCP packet is broadcast out from the interface. Use the no operator to delete the forwarding path from the client to the server: `no ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D>`.
fwd-path <A.B.C.D> <A.B.C.D> disable	Disables DHCP relaying on the path from the IP address to the server. This feature is disabled by default. A.B.C.D is the IP address configured on an interface (a locally configured IP address). A.B.C.D is the IP address of the DHCP server in the network.
fwd-path <A.B.C.D> <A.B.C.D> enable	Enables DHCP relaying on the path from the IP address to the server. A.B.C.D is the IP address configured on an interface (a locally configured IP address). A.B.C.D is the IP address of the DHCP server in the network. If this IP address corresponds to the locally configured IP network, the DHCP packet is broadcast out from the interface. Use the no operator to disable DHCP: `no dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> enable`. To configure this option to the default value, use the default operator with the `ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D>` command.

Variable	Value
fwd-path <A.B.C.D> <A.B.C.D> mode <bootp|bootp_dhcp|dhcp>	Modifies DHCP mode to forward BootP messages only, DHCP messages only, or both. The default is both. mode is {bootp | bootp_dhcp | dhcp}.

Showing DHCP relay information

Display relay information to show relay information about DHCP routes and counters.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display information about DHCP relay forward paths:

   ```
   show ip dhcp-relay fwd-path [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

3. Display information about DHCP relay counters:

   ```
   show ip dhcp-relay counters [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

4. Display the options for each listed interface:

   ```
   show ip dhcp-relay interface [gigabitethernet {slow/port [slot/port]
   [,...]}] [vlan <1-4084>] [vrf WORD <1-16>] [vrfids WORD <0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
VSP-9012:1(config)#show ip dhcp-relay interface
================================================================================
                                  Port Dhcp
================================================================================
PORT    VRF              MAX MIN          ALWAYS CIRCUIT    REMOTE TRUST
NUM     NAME          ENABLE  HOP SEC      MODE  BCAST  ID        ID     CIRC
--------------------------------------------------------------------------------
4/3     GlobalRouter false   4    0       both  false


================================================================================
                                  Vlan Dhcp
================================================================================
VLAN VRF                     MAX MIN               ALWAYS CIRCUIT REMOTE TRUST
ID   NAME             ENABLE HOP SEC   MODE         BCAST  ID      ID     CIRC
--------------------------------------------------------------------------------

All 0 out of 0 of Vlan Dhcp Entries displayed
```

Variable definitions

Use the data in the following table to use the **show ip dhcp-relay** command.

Table 17: Variable definitions

Variable	Value
vrf WORD<1-16>	The name of the VRF.
vrfids WORD<0-512>	The ID of the VRF. The value is an integer in the range of 0–512.

Use the data in the following table to use the **show ip dhcp-relay interface** command.

Variable	Value
[gigabitethernet {slot/port[-slot/port][,...]}]	Specifies the slot and port or range of slots and ports for the Gigabit Ethernet interface type.
[vlan <1-4084>]	Specifies the VLAN id in the range of 1 to 4084.
[vrf WORD<1-16>]	Specifies the name of the VRF.
[vrfids WORD<0-512>]	Specifies the ID of the VRF. The value is an integer from 0– 512.

Configuring DHCP option 82

Configure the DHCP option 82 to enable the circuit ID to encode an agent-local identifier of the circuit from which a DHCP client-to-server packet is received. Configure the DHCP option 82 to enable the remote ID to encode the mac address of the interface on which the packet is received. By default, the DHCP option 82 is disabled.

Before you begin

- To configure the DHCP option 82 on a VLAN, you must enter the VLAN Interface Configuration mode.
- To configure the DHCP option 82 on a brouter port, you must enter the GigabitEthernet Interface Configuration mode.
- You must enable ip and dhcp-relay on the VLAN.

Procedure

1. Enable the circuit ID:

   ```
   ip dhcp-relay circuitID
   ```

2. Enable the remote ID:

   ```
   ip dhcp-relay remoteID
   ```

3. Configure the circuit as trusted:

   ```
   ip dhcp-relay trusted
   ```

4. Show statistics for option 82, which is the relay agent information option:

   ```
   show ip dhcp-relay counters option82 [vrf WORD <1-16>] [vrfids WORD
   <0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 4/10
```

Enable the circuit ID:

```
VSP-9012:1(config-if)#ip dhcp—relay circuitID
```

Enable the remote ID:

```
VSP-9012:1(config-if)#ip dhcp-relay remoteID
```

Configure the circuit as trusted:

```
VSP-9012:1(config-if)#ip dhcp-relay trusted
```

Show statistics for option 82, which is the relay agent information option:

```
VSP-9012:1(config-if)#show ip dhcp-relay counters option82
```

Variable definitions

Use the data in the following table to configure the DHCP option 82 through ACLI.

Table 18: Variable definitions

Variable	Value
circuitID	Enables the Circuit ID.
remoteID	Enables the Remote ID.
trusted	Sets the circuit as trusted.

Use the data in the following table to use the **`show ip dhcp-relay counters option82 [vrf WORD <1–16>] [vrfids WORD <0–512>]`** command.

Variable	Value
vrf WORD <1–16>	Displays DHCP counters for a particular VRF. WORD <1–16> specifies the VRF name.
vrfids WORD <0–512>	Displays a DHCP forward path for a particular VRF. WORD <0–512> specifies the VRF ID.

Configuring DHCP relay on a port or VLAN

You can view and configure the DHCP parameters on specific ports or on a VLAN.

Before you begin

- You must configure IP on the interface.

Procedure

1. Enter Interface Configuration mode:

   ```
   enable
   configure terminal
   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```
   or `interface vlan <1-4084>`

2. Enable DHCP parameters on a specified port or VLAN:

   ```
   ip dhcp-relay
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 4/10
```

Enable DHCP parameters on a specified port or VLAN:

```
VSP-9012:1(config-if)#ip dhcp-relay
```

Variable definitions

Use the data in the following table to use the **`ip dhcp-relay`** command.

Use the **`no`** operator to disable DHCP parameters on specified ports: **`no ip dhcp-relay`**.

Note:

The **`no ip dhcp-relay`** command disables DHCP Relay, it does not delete the DHCP entry.

To configure this option to the default value, use the **`default`** operator with this command.

Table 19: Variable definitions

Variable	Value
broadcast	Enables the device to send the server reply as a broadcast to the end station. After you disable this variable, the device sends the server reply as a unicast to the end station. Use the no operator to disable broadcast: `no ip dhcp-relay broadcast`.

Variable	Value
	To configure this option to the default value, use the default operator with this command.
circuitId	Enables the device to insert the Option 82 Circuit ID into the packets sent to the server (enables DHCP Option 82).
fwd-path <A.B.C.D> [vrid <1-255>]	Creates a forward path server with a virtual router ID (or VRRP ID), a mode, and a state. A.B.C.D is the IP address. vrid <1-255> is the ID of the virtual router and is an integer from 1 to 255. Use the no operator to delete a forward path server with a specific value and virtual router ID: `no ip dhcp-relay fwd-path <A.B.C.D> [vrid <1-255>]` To configure this option to the default value, use the default operator with this command.
fwd-path <A.B.C.D> disable [vrid <1-255>]	Disables a forward path server with a specific value and virtual router ID. A.B.C.D is the IP address. vrid <1-255> is the ID of the virtual router (or VRRP ID) and is an integer from 1 to 255.
fwd-path <A.B.C.D> enable [vrid <1-255>]	Enables a forward path server with a specific value and virtual router ID (or VRRP ID). A.B.C.D is the IP address in the form a.b.c.d. vrid <1-255> is the ID of the virtual router and is an integer from 1 to 255.
fwd-path <A.B.C.D> mode <bootp|bootp_dhcp|dhcp> [vrid <1-255>]	Configures the forward path mode for a VLAN. This command string is available only in VLAN Interface Configuration mode. A.B.C.D is the IP address in the form a.b.c.d. mode is a choice of bootp, dhcp, or bootp_dhcp. vrid <1-255> is the ID of the virtual router (or VRRP ID) and is an integer from 1 to 255. To configure this option to the default value, use the default operator with this command.
max-hop <1-16>	Configures the maximum number of hops before a BootP/DHCP packet is discarded (1 to 16). The default is 4. To configure this option to the default value, use the default operator with this command.
min-sec <0-65535>	Configures the minimum seconds count for DHCP. If the secs field in the BootP/DHCP packet header is greater than this value, the device relays or forwards the packet; otherwise, the packet is dropped (0 to 65535). The default is 0 seconds.

Variable	Value
	To configure this option to the default value, use the default operator with this command.
mode <bootp|bootp_dhcp|dhcp>	Configures DHCP mode to forward BootP messages only, DHCP messages only, or both. The default is both. To configure this option to the default value, use the default operator with this command.
remoteId	Enables the device to insert the Option 82 Remote ID into packets sent to the server (enables DHCP Option 82).
trusted	Configures the circuit as trusted in an Option 82 context.

Configuring UDP broadcast forwarding

About this task

By default, routers do not forward broadcasts. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts. You must set up UDP broadcast forwarding on the system. Configure UDP broadcast forwarding to forward the UDP broadcasts of network applications to the required server through physical or virtual router interfaces.

Procedure

1. Enter protocols into a table.
2. Create policies (protocol/server pairs).
3. Assemble the policies into lists or profiles.
4. Apply the list to the appropriate interfaces.

Configuring UDP protocols

Configure UDP protocols to determine which UDP broadcasts are forwarded.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a UDP protocol:

   ```
   ip forward-protocol udp <1-65535> WORD<1-15>
   ```

3. Confirm your configuration:

```
show ip forward-protocol udp [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#ip forward-protocol udp 53 DNS
```

Confirm your configuration:

```
show ip forward-protocol udp
```

Variable definitions

Use the data in the following table to use the **`ip forward-protocol udp`** command.

Table 20: Variable definitions

Variable	Value
<1-65535> WORD<1-15>	Creates a new UDP protocol. <1-65535> WORD<1-15> is the UDP protocol name as a string. Use the no operator to delete a UDP protocol `no ip forward-protocol udp <1-65535>`.
[vrf WORD<1-16>]	Specifies the name of the VRF.
[vrfids WORD<0-512>]	Specifies the ID of the VRF. The value is an integer from 0–512.

Configuring a UDP port forward entry

Configure a UDP port forward entry to add or remove a port forward entry.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a UDP port forward entry:

   ```
   ip forward-protocol udp portfwd <1-65535> {A.B.C.D}
   ```

3. Confirm your configuration:

   ```
   show ip forward-protocol udp portfwd [vrf WORD<1-16>] [vrfids
   WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Configure a UDP port forward entry:

```
VSP-9012:1(config)#ip forward-protocol udp portfwd 150 30.30.1.1
```

Variable definitions

Use the data in the following table to use the **`ip forward-protocol udp portfwd`** command.

Table 21: Variable definitions

Variable	Value
<1-65535> {A.B.C.D}	Adds a UDP protocol port to the specified port forwarding list. *1-65535* is a UDP protocol port in the range of 1–65535. A.B.C.D is an IP address in a.b.c.d format. Use the no operator to remove a protocol port forwarding entry and IP address from the list: `no ip forward-protocol udp portfwd <1-65535> <A.B.C.D>`. To configure this option to the default value, use the default operator with this command.
[vrf WORD<1-16>]	Specifies the name of the VRF.
[vrfids WORD<0-512>]	Specifies the ID of VRF and is an integer from 0–512.

Configuring the UDP port forwarding list

Before you begin

- You must log on to the Global Configuration mode, the VLAN Interface Configuration mode, or the VRF Router Configuration mode in ACLI.

About this task

Configure the UDP port forwarding list to assign protocols and servers to the port forward list.

Procedure

1. Configure the UDP port forwarding list:

   ```
   ip forward-protocol udp portfwdlist <1-1000>
   ```

Important:

The following two steps are not available in the Global Configuration or VRF Router Configuration mode. The following two commands are available in VLAN Interface Configuration mode only.

2. Log on to Interface Configuration mode:

   interface vlan

3. Configure the broadcast mask:

   ```
   ip forward-protocol udp broadcastmask {A.B.C.D}
   ```

4. Configure the maximum time to live:

   ```
   ip forward-protocol udp maxttl <1-16>
   ```

5. Confirm your configuration:

   ```
   show ip forward-protocol udp portfwdlist [vrf WORD<1-16>] [vrfids
   WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Configure the UDP port forwarding list:

```
VSP-9012:1(config)#ip forward-list protocol udp portfwdlist 1
```

Log on to the VLAN interface configuration mode:

```
VSP-9012:1(config)# interface vlan 3
```

Configure the broadcast mask:

```
VSP-9012:1(config-if)#ip forward-protocol udp broadcastmask
100.31.255.255
```

Configure the maximum time to live:

```
VSP-9012:1(config-if)#ip forward-protocol udp maxttl 10
```

Confirm the configuration:

```
VSP-9012:1(config-if)#show ip forward-protocol udp portfwdlist
```

Variable definitions

Use the data in the following table to use the **`ip forward-protocol udp`** command.

Table 22: Variable definitions

Variable	Value
<1-1000>	Creates a UDP port forwarding list in the range of 1–1000.
<1–65535> {A.B.C.D}	Adds a UDP protocol port to the specified port forwarding list. *1-65535* is a UDP protocol port in the range of 1–65535. A.B.C.D is an IP address in a.b.c.d format. Use the no operator to remove or delete a port forwarding list ID, `no ip forward-protocol udp portfwdlist <1-1000> <1-65535> <A.B.C.D>`. To configure this option to use the default value, use the default operator with this command.
name WORD<0–15>	Changes the name of the port forwarding list.

Use the data in the following table to use the **`ip forward-protocol udp`** command.

Variable	Value
broadcastmask {A.B.C.D}	Configures the interface broadcast mask (the interface broadcast mask can be different from the interface mask). A.B.C.D is an IP address in a.b.c.d format. Use the no operator to delete the broadcast mask: `no ip forward-protocol udp broadcastmask {A.B.C.D}` To configure this option to the default value, use the default operator with this command.
maxttl <1-16>	Configures the maximum time-to-live value (TTL) for the UDP broadcast forwarded by the interface. The range is 1–16.
portfwdlist <1–1000>	Assigns the list to the VLAN.
vlan <1–4084> [portfwdlist <1–1000>]	Specifies the VLAN ID. If you use the portfwdlist variable with the vlan variable, it assigns the list to the specified VLAN, regardless of which VLAN context you currently configure.

Showing UDP forward information

Before you begin

- You must log on to Privileged EXEC mode, Global Configuration mode, or the VRF Router Configuration mode in ACLI.

About this task

Show UDP forward information to view information about the UDP forwarding characteristics of the device. Four show options exist:

- Show the interface information
- Show the port forward information
- Show the port forward list information
- Show the protocol information

Procedure

1. Display information about the UDP interface for all IP addresses or a specified IP address:

   ```
   show ip forward-protocol udp interface [<A.B.C.D>] [vrf WORD<1-16>]
   [vrfids WORD<0-512>]
   ```

2. Display the UDP port forwarding table:

   ```
   show ip forward-protocol udp portfwd [vrf WORD<1-16>] [vrfids
   WORD<0-512>]
   ```

3. Display the UDP port forwarding list table for the specified list or all lists on the device:

   ```
   show ip forward-protocol udp portfwdlist [vrf WORD<1-16>] [vrfids
   WORD<0-512>]
   ```

4. Display the UDP protocol table with the UDP port numbers for each supported or designated protocol:

   ```
   show ip forward-protocol udp [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#show ip forward-protocol udp
================================================================================
                        Udp Protocol Tbl - GlobalRouter
================================================================================
UDP_PORT PROTOCOL_NAME
--------------------------------------------------------------------------------
37       Time Service
49       TACACS Service
53       DNS
69       TFTP
137      NetBIOS NameSrv
138      NetBIOS DataSrv
5050     test50
```


Comments? infodev@avaya.com

Variable definitions

Use the data in the following table to use the **`show ip forward-protocol udp interface`** command.

Table 23: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the IP address for the interface in a.b.c.d format.
vrf *WORD<1–16>*	Specifies the name of the VRF.
vrfids *WORD<0–512>*	Specifies the ID of the VRF and is an integer in the range of 0 to 512.

Chapter 7: DHCP and UDP configuration using Enterprise Device Manager

Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), dynamically provides host configuration information to workstations. To lower administrative overhead, network managers prefer to configure a small number of DHCP servers in a central location. Using few DHCP servers requires the routers connecting to the subnets or bridge (or VLAN) domains to support the BootP/DHCP relay function so that hosts can retrieve the configuration information from servers several router hops away.

User datagram protocol (UDP) is a connectionless protocol that adds reliability and multiplexing to IP. It describes how messages reach application programs within a destination computer. Some network applications, such as the NetBIOS name service, rely on a UDP broadcast to request a service or to locate a service. By default, broadcasts are not forwarded by a router. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts.

Important:

BootP/DHCP relays are supported only on IP routed port-based VLANs and protocol-based VLANs. BootP/DHCP relays are not supported on IP subnet-based VLANs.

Before you begin

You must enable DHCP relay on the path for port or VLAN configuration to take effect.

Configuring DHCP on a brouter port or a VRF instance

Before you begin

- You must first enable BootP/DHCP relay on a port (or VLAN).
- You must enable DHCP and forwarding path.
- You must enable IP Routing on the interface.

About this task

Use the DHCP tab to configure the DHCP behavior on a brouter port or a VRF instance. The DHCP tab is available only if the port is routed (that is, assigned an IP address).

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **DHCP Relay** tab.
5. Click **Enable** to select the DHCP option. The default is disable.
6. Configure the other parameters as needed.
7. Click **Apply**.

DHCP field descriptions

Use the data in the following table to use the **DHCP** tab.

Name	Description
Enable	Lets you use BootP/DHCP on the port. The default is disable.
MaxHop	Sets the maximum number of hops before a BootP/DHCP packet is discarded (1 to 16). The default is 4.
MinSec	The secs field in the BootP/DHCP packet header represents the elapsed time since the client first sent the message. If the secs field in the packet header is greater than this value, the system relays or forwards the packet; otherwise, the packet is dropped. The default is 0 seconds.
Mode	Sets the interface to process only BootP, only DHCP, or both types of packets. The default is both.
AlwaysBroadcast	When enabled, the server reply is sent as a broadcast back to the end station. The default is disable.
CircuitID	When enabled, the VSP DHCP Relay inserts the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
RemoteID	When enabled, the VSP DHCP Relay inserts the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Trusted	When enabled, the DHCP server receives the DHCP packets through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default is disable.

Configuring BootP/DHCP on a VLAN or VRF instance

Before you begin

- You must enable IP Routing on the interface.

About this task

Use the DHCP Relay tab to configure the DHCP behavior on a VLAN. The DHCP Relay tab is available only if the VLAN is routed and is assigned an IP address.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP**.
5. Click the **DHCP Relay** tab.
6. Select **Enable**.
7. Configure the parameters as required.
8. Click **Apply**.

DHCP Relay field descriptions

Use the data in the following table to use the **DHCP Relay** tab.

Variable	Value
Enable	Lets you use BootP/DHCP on the port. The default is disable.
MaxHop	Sets the maximum number of hops a BootP/DHCP packet can take from the DHCP client to the DHCP server. The maximum number of hops is 16. The default is 4.
MinSec	Represents the minimum number of seconds to wait between receiving a DHCP packet and forwarding the DHCP packet to the DHCP server. A value of 0 indicates that forwarding is done immediately. The default value is 0.
Mode	Indicates the type of DHCP packet required. The options are: • bootp • dhcp • both

Variable	Value
	The default is both.
AlwaysBroadcast	When enabled, the DHCP Reply packets are sent as a broadcast to the DHCP client. The default is disable.
CircuitID	When enabled, the VSP DHCP Relay inserts the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
RemoteID	When enabled, the VSP DHCP Relay inserts the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Trusted	When enabled, the DHCP server receives the DHCP packets through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default is disable.

Configuring DHCP relay

About this task

After you configure the BootP/DHCP relay on an IP interface, you can configure forwarding paths to indicate where packets are forwarded. The forwarding paths are based on the type of packet and where the packet is received.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **DHCP Relay**.
3. Click **Insert**.
4. In the **AgentAddr** box, type the agent address.
5. In the **ServerAddr** list, type the server address.
6. Click **Enable** to enable BootP/DHCP relay. You can enable or disable each agent server forwarding policy. The default is enabled.
7. In the **Mode** box, select the type of messages to relay.

   Both the mode setting for the DHCP interface and the mode setting for the agent interface determine which packets are forwarded.
8. Click **Insert**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AgentAddr	The IP address of the input interface (agent) on which the BootP/DHCP request packets are received for forwarding. This address is the IP address of either a brouter port or a VLAN for which forwarding is enabled.
ServerAddr	This parameter is either the IP address of the BootP/DHCP server or the address of another local interface. • If it is the address of the BootP/DHCP server, the request is unicast to the server address. • If the address is one of the IP addresses of an interface on the system, the BootP/DHCP requests are broadcast out of that local interface.
Enable	Enables BootP/DHCP relay.
Mode	Specifies the type of messages relayed: • Only BootP • Only DHCP • Both types of messages The default is to forward both BootP and DHCP messages.

Viewing DHCP relay configuration information

About this task

Use the DHCP Relay Interfaces tab to view configuration information about the DHCP relay. To change the configuration information, double-click the value in the field under the required interface, and enter a new value.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **DHCP Relay**.
3. Click the **Interfaces** tab.

Interfaces field descriptions

Use the data in the following table to use the **Interfaces** tab.

Variable	Value
IfIndex	A read-only interface number that represents a physical interface, or the VLAN logical interface.
MaxHop	Sets the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The maximum number of hops is 16. The default is 4.
MinSec	Represents the minimum number of seconds to wait between receiving a DHCP packet and forwarding the DHCP packet to the DHCP server. A value of 0 indicates that forwarding is done immediately. The default value is 0.
Mode	Indicates the type of DHCP packet required. The options are: • bootp • dhcp • both The default is both.
AlwaysBroadcast	Indicates if DHCP Reply packets can be sent as a broadcast to the DHCP client. The default is false.
CircuitId	Indicates if the VSP DHCP Relay inserted the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
RemoteId	Indicates if the VSP DHCP Relay inserted the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Trusted	Indicates if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.

Managing UDP forwarding protocols

The Avaya Virtual Services Platform 9000 configures the following protocols, by default:

- Time Service
- Terminal Access Controller Access Control System (TACACS) Service
- DNS
- Trivial file transfer protocol (TFTP)
- Network Basic Input/Output System (NetBIOS) NameSrv

- NetBIOS DataSrv

You can use these protocols to create forwarding entries and lists but you cannot delete them; you can add or remove other protocols to the list of protocols.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **UDP Forwarding**.
3. Click the **Protocols** tab.
4. Click **Insert**.
5. In the **PortNumber** field, type a UDP port number.

   This number defines the UDP port used by the server process as its contact port. The range is from 1 to 65535 and cannot be one of the UDP port numbers or a number previously assigned.
6. In the **Name** field, type a name for the protocol.
7. Click **Insert**.

   The protocol is added to the Protocol table. After you create a protocol, you cannot change its name or number.

Protocols field descriptions

Use the data in the following table to use the **Protocols** tab.

Name	Description
PortNumber	Defines the UDP port (1 to 65535).
Name	Specifies an administratively assigned name for this list (0 to 15 characters).

Managing UDP forwarding

About this task

You manage UDP forwarding by defining the destination addresses for the UDP protocol.

Procedure

1. In the navigation tree, expand the following folders:**Configuration** > **IP**.
2. Click **UDP Forwarding**.
3. Click the **Forwardings** tab.
4. Click **Insert**.

5. In the Insert Forwardings dialog box, select a destination UDP port from the defined protocols in the **DestPort** box.
6. Enter a destination IP address in the **DestAddr** box.

   The destination address can be any IP server address for the protocol application or the IP address of an interface on the router.
7. Click **Insert**. The information is added to the Forwarding tab.

Forwardings field descriptions

Use the data in the following table to use the **Forwardings** tab.

Name	Description
DestPort	Specifies the port number defined for UDP, depending upon the protocol type.
DestAddr	Specifies the destination address can be any IP server address for the protocol application or the IP address of an interface on the router: • If the address is that of a server, the packet is sent as a unicast packet to this address. • If the address is that of an interface on the router, the frame is rebroadcast.
Id	Specifies an integer that identifies this entry internally.
NumFwdPackets	Specifies the total number of UDP broadcast packets forwarded using this policy.
NumDropPacketsTtlExpired	Specifies the total number of UDP broadcast packets dropped because the time-to-live value (TTL) expired.
NumDropPacketsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the specified destination address was unreachable.

Creating the forwarding profile

About this task

A forwarding profile is a collection of port and destination pairs. When you configure UDP forwarding list entries, be sure to first configure the UDP forwarding list. Then, configure your UDP forwarding list entries and assign them to a UDP forwarding list. If you do not assign a UDP forwarding list entry to at least one UDP forwarding list, the UDP forwarding list is lost after a restart.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.

Comments? infodev@avaya.com

2. Click **UDP Forwarding**.
3. Click the **Forwarding Lists** tab.
4. Click **Insert**.
5. In the **Id** field, type the forwarding list ID.
6. In the **Name** field, type the name of the forwarding list if required.

   The forwarding list appears in the **FwdIdList** box.
7. Click **Insert**.

Forwarding Lists field descriptions

Use the data in the following table to use the **Forwarding Lists** tab and **Insert Forwarding Lists** dialog box.

Name	Description
Id	Specifies a value that uniquely identifies this list of entries (1 to 1000).
Name	Specifies an administratively assigned name for this list (0 to 15 characters).
FwdIdList	Specifies the zero or more port forwarding entries associated with this list. Each list identifier is stored as 2 bytes in this array, starting from 0 bytes (size=64). Clicking on the ellipsis (...) button in this field displays the ID list.

Managing the broadcast interface

About this task

Manage the broadcast interface by specifying and displaying which router interfaces can receive UDP broadcasts to forward.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **UDP Forwarding**.
3. Click the **Broadcast Interfaces** tab.
4. Click **Insert**.
5. In the **LocalIfAddr** field, click the ellipsis **(...)** to select a local interface IP address from the list, and then click **OK**.
6. In the **UdpPortFwdListId** field, click the ellipsis **(...)** to select a forwarding list ID from the list, and then click **OK**.
7. In the **MaxTtl** field, type the maximum number of hops an IP broadcast can take from the source device to the destination device (the default is 4; the range is 1 to 16).

8. In the **BroadCastMask** field, enter the subnet mask of the local interface that broadcasts the UDP broadcast packets.

   When you configure the UDP forwarding broadcast mask, the broadcast mask must be less specific (shorter in length) or equally specific (equal in length) to the subnet mask of the IP interface on which it is configured. If the UDP forwarding broadcast mask is more specific than the subnet mask of the corresponding IP interface, UDP forwarding does not function properly.

9. Click **Insert**.

Broadcast Interfaces field descriptions

Use the data in the following table to use the **Broadcast Interfaces** tab.

Name	Description
LocalIfAddr	Specifies the IP address of the local router interface that receives forwarded UDP broadcast packets.
UdpPortFwdListId	Specifies the number of the UDP lists or profiles that this interface is configured to forward (0 to100). A value of 0 indicates that the interface cannot forward any UDP broadcast packets.
MaxTtl	Specifies the maximum number of hops an IP broadcast packet can take from the source device to the destination device (the default is 4; the range is 1 to 16).
NumRxPkts	Specifies the total number of UDP broadcast packets received by this local interface.
NumFwdPkts	Specifies the total number of UDP broadcast packets forwarded by this local interface.
NumDropPktsMaxTtlExpired	Specifies the total number of UDP broadcast packets dropped because the time-to-live (TTL) value expired.
NumDropPktsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the destination was unreachable.
NumDropPktsUnknownPort	Specifies the total number of UDP broadcast packets dropped because the destination port or protocol specified has no matching forwarding policy.
BroadCastMask	Specifies the subnet mask of the local interface that broadcasts the UDP broadcast packets.

Viewing UDP endpoint information

View UDP Endpoints to confirm correct configuration.

About this task

You can use UDP endpoint information to display local and remote UDP activity.

Since UDP is a protocol used to establish connectionless network sessions, you need to monitor local and remote UDP activity and to know which applications are running over UDP.

You can determine which applications are active by checking the port number.

Processes are further identified with a UDP session to allow for the multiplexing of a port mapping for UDP.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP** or **Configuration** > **IPv6**.
2. Click **TCP/UDP**.
3. Click the **UDP Endpoints** tab.

UDP Endpoints field descriptions

Use the data in the following table to use the **UDP Endpoints** tab.

Name	Description
LocalAddressType	Displays the local address type (IPv6 or IPv4).
LocalAddress	Displays the local IPv6 address.
LocalPort	Displays the local port number.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IPv6 address.
RemotePort	Displays the remote port number.
Instance	Distinguishes between multiple processes connected to the UDP endpoint.
Process	Displays the ID for the UDP process.

Chapter 8: IP policy configuration using ACLI

Configure IP policies to form a unified database of route policies that Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) can use for filtering tasks.

A policy is identified by a name or an ID. Under a given policy you can have several sequence numbers, each of which is equal to one policy in the old convention. Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference field, use only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce and redistribute purposes.

You can apply one policy for one purpose, for example, RIP announce on a RIP interface. All sequence numbers under the given policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

Configuring prefix lists

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

About this task

Important:

When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a prefix list:

   ```
   ip prefix-list WORD<1-64> {A.B.C.D/X} [ge <0-32>] [le <0-32>]
   ```

3. **(Optional)** Rename an existing prefix list:

   ```
   ip prefix-list WORD<1-64> name WORD<1-64>
   ```

4. Display the prefix list:

   ```
   show ip prefix-list [prefix {A.B.C.D}] [vrf WORD<1-16>] [vrfids
   WORD<0-512>] [WORD <1-64>]
   ```

Example

Configure a prefix-list. Display the prefix list.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#ip prefix-list LIST1 47.17.121.50/255.255.255.0
VSP-9012:1(config)#show ip prefix-list LIST1
================================================================================
                          Prefix List - GlobalRouter
================================================================================

        PREFIX           MASKLEN FROM TO
--------------------------------------------------------------------------------

List 1    LIST1:
        47.17.121.50     24      24   24
1 Total Prefix List entries configured
--------------------------------------------------------------------------------
Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
```

Variable definitions

Use the data in the following table to use the **`ip prefix-list`** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats: • a.b.c.d/x • a.b.c.d/x.x.x.x • default
ge <0–32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
le <0–32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.

Variable	Value
name WORD<1-64>	Renames the specified prefix list. The name length is 1–64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

Use the data in the following table to use the **`show ip prefix-list`** command.

Variable	Value
{A.B.C.D}	Specifies the prefix to include in the command output.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Use the following table to use the **`show ip prefix-list`** command output.

Variable	Value
PREFIX	Indicates the member of a specific prefix list.
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
TO	Indicates the prefix mask endpoint in bits.

Configuring IP route policies

Configure a route policy so that the device can control routes that certain packets can take. For example, you can use a route policy to deny certain Border Gateway Protocol (BGP) routes.

The route policy defines the matching criteria and the actions taken if the policy matches.

About this task

After you create and enable the policy, you can apply it to an interface. You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In this case, all sequence numbers under the given policy apply to that filter.

Create and enable the policy for IS-IS accept policies for Avaya VENA Fabric Connect for Layer 3 Virtual Services Networks (VSNs) and IP Shortcuts, then apply the IS-IS accept policy filters. For more information on IS-IS accept policy filters, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Note:

After you configure route-map in Global Configuration mode or VRF Router Configuration mode, the device enters Route-Map Configuration mode, where you configure the action the policy takes, and define other fields the policy enforces.

Procedure

1. Enter Route-Map Configuration mode:

   ```
   enable
   configure terminal
   route-map WORD<1-64> <1-65535>
   ```

2. At the route-map prompt, define the fields the policy enforces:

   ```
   match metric <0-65535>
   ```

   In this procedure, the metric field is used. You can configure more than one field.

3. Define the action the policy takes to allow the route:

   ```
   permit
   ```

4. Define the action the policy takes to ignore the route:

   ```
   no permit
   ```

5. Configure other policy parameters as required. Use the following variable definitions table for other parameters.

6. Display current information about the IP route policy:

   ```
   show route-map [WORD<1-64>] [<1-65535>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

Enter Route-Map Configuration mode. At the route-map prompt, define the fields the policy enforces. Define the action the policy takes. Display current information about the IP route policy.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#route-map RedisStatic 1
Switch:1(route-map)# match metric 0
Switch:1(route-map)# permit
Switch:1(route-map)# show route-map RedisStatic
================================================================================
                          Route Policy - GlobalRouter
================================================================================

NAME                                                            SEQ   MODE EN
--------------------------------------------------------------------------------
RedisStatic                                                     1     PRMT DIS
```

Variable definitions

Use the data in the following table to use the **match** command.

Table 24: Variable definitions

Variable	Value
as-path WORD<0-256>	Configures the device to match 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 <0-256>* specifies the list IDs of up to four AS-lists, separated by a comma. Use the no operator to disable match as-path: `no match as-path WORD<0-256>`
community WORD<0-256>	Configures the device to match 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 <0-256>* specifies the list IDs of up to four defined community lists, separated by a comma. Use the no operator to disable match community: `no match community WORD<0-256>`
community-exact enable	When disabled, configures the device so match community-exact results in a match when the community attribute of the BGP routes match an entry of a community-list specified in match-community. When enabled, configures the device so 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. enable enables match community-exact. Use the no operator to disable match community-exact: `no match community-exact enable`
interface WORD <1–64>	If configured, configures the device to match 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 route types. *WORD <1–64>* specifies the name of up to four defined prefix lists, separated by a comma. Use the no operator to disable match-interface: `no match interface WORD <1-64>`
metric <0-65535>	Configures the device to match the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored. *<0-65535>* specifies the metric value. The default is 0.
network WORD <1–64>	Configures the device to match the destination network against the contents of the specified prefix lists.

Variable	Value
	WORD <1–64> specifies the name of up to four defined prefix lists, separated by a comma. Use the no operator to disable match network: `no match network WORD <1-64>`.
next-hop WORD<1–64>	Configures the device to match the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes. *WORD <1–64>* specifies the name of up to four defined prefix lists, separated by a comma. Use the no operator to disable match next hop: `no match next-hop WORD<1-64>`.
protocol WORD<0-60>	Configures the device to match the protocol through which the route is learned. *WORD <0-60>* is |xxx, where xxx is local, ospf, ebgp, ibgp, isis, rip, static, or a combination separated by |, in a string length 0–60. Use the no operator to disable match protocol: `no match protocol WORD<0-60>`
route-source WORD<1–64>	Configures the system to match the next-hop 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. *WORD <1–64>* specifies the name of up to four defined prefix lists, separated by a comma. Use the no operator to disable match route source: `no match route-source WORD<1-64>`
route-type {any|local|internal|external|external-1|external-2}	Configures a specific route type to match (applies only to OSPF routes). any|local|internal|external|external-1|external-2 specifies OSPF routes of the specified type only (External-1 or External-2). Another value is ignored.
tag WORD<0-256>	Specifies a list of tags used during the match criteria process. Contains one or more tag values. *WORD<0-256>* is a value from 0–256.
[vrf WORD<1-16>] [vrfids WORD<0-512>]	Configures a specific VRF to match (applies only to RIP routes).
set community-mode <additive|none|unchanged>	Configures the community mode. additive—the device prepends the community number of the community list specified in set-community to the old

Variable	Value
	community path attribute of the BGP routes that match this policy. none—the device removes the community path attribute of the BGP routes that match this policy to the specified value.

Use the data in the following table to use the **set** command.

Table 25: Variable definitions

Variable	Value
as-path WORD<0-256>	Configures the device to add the AS number of the AS-list to the BGP routes that match this policy. *WORD<0-256>* specifies the list ID of up to four defined AS-lists separated by a comma. Use the no operator to delete the AS number: `no set as-path WORD<0–256>`
as-path-mode <tag|prepend>	Configures the AS path mode. Prepend is the default configuration. The device prepends the AS number of the AS-list specified in set-as-path to the old as-path attribute of the BGP routes that match this policy.
automatic-tag enable	Configures the tag automatically. Used for BGP routes only. Use the no operator to disable the tag: `no set automatic-tag enable`
community WORD<0-256>	Configures the device to add the community number of the community list to the BGP routes that match this policy. *WORD <0-256>* specifies the list ID of up to four defined community lists separated by a comma. Use the no operator to delete the community number: `no set community WORD<0–256>`
community-mode <additive|none|unchanged>	Configures the community mode. additive—the device 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 device removes the community path attribute of the BGP routes that match this policy to the specified value.
injectlist WORD<1–64>	Configures the device to replace the destination network of the route that matches this policy with the contents of the specified prefix list. *WORD<1–64>* specifies one prefix list by name.

Variable	Value
	Use the no operator to disable set injectlist: `no set injectlist`
ip-preference <0-255>	Configures the preference. This applies to accept policies only. *<0-255>* is the range you can assign to the routes.
local-preference <0-65535>	Configures the device to match the local preference, applicable to all protocols. *<0–655356>* specifies the preference value.
mask <A.B.C.D>	Configures the mask of the route that matches this policy. This applies only to RIP accept policies. *A.B.C.D* is a valid contiguous IP mask. Use the no operator to disable set mask: `no set mask`
metric <0-65535>	Configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF for RIP, the original cost of the route or default-import-metric is used (applies to IS-IS routes also).
metric-type {type1|type2}	Configures the metric type for the routes to announce into the OSPF domain that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.
next-hop <A.B.C.D>	Specifies the IP address of the next-hop router. Use the no operator to disable set next-hop: `no set next-hop`
nssa-pbit enable	Configures the not so stubby area (NSSA) translation P bit. Applicable to OSPF announce policies only. Use the no operator to disable set nssa-pbit: `no set nssa-pbit enable`
origin {igp|egp|incomplete}	Configures the device to change the origin path attribute of the BGP routes that match this policy to the specified value.
origin-egp-as <0-65535>	Indicates the remote autonomous system number. Applicable to BGP only.
tag <0-65535>	Configures the tag of the destination routing protocol. If not specified, the device forwards the tag value in the source routing protocol. A value of 0 indicates that this parameter is not configured.
weight <0-65535>	Configures the weight value for the routing table. For BGP, this value overrides the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. Used for BGP only. A value of 0 indicates that this parameter is not configured.

Use the data in the following table to use the **name** command.

Table 26: Variable definitions

Variable	Value
WORD<1-64>	Renames a policy and changes the name field for all sequence numbers under the given policy.

Job aid

Use the data in the following table to use the `show route-map` command output.

Table 27: Variable definitions

Variable	Value
NAME	Indicates the name of the route policy.
SEQ	Indicates the second index used to identify a specific policy within the route policy group (grouped by ID). Use this field to specify different match and set parameters and an action.
MODE	Indicates the action to take when this policy is selected for a specific route. Options are permit, deny, or continue. Permit indicates to allow the route. Deny indicates to ignore the route. Continue means continue checking the next match criteria configured in the next policy sequence; if none, take the default action in the given context.
EN	Indicates whether this policy is enabled. If disabled, the policy is not used.

Configuring a policy to accept external routes from a router

Perform this procedure to configure a policy to accept external routes from a specified advertising router.

For more information on IS-IS accept policy filters for Avaya VENA Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

1. Enter OSPF Router Configuration mode:

   ```
   enable
   configure terminal
   ```

```
router ospf
```

2. Create a policy to accept external routes from a specified advertising route:

   ```
   accept adv-rtr <A.B.C.D>
   ```

3. Exit to the Privileged EXEC mode.
4. Apply the OSPF accept policy change:

   ```
   ip ospf apply accept adv-rtr <A.B.C.D>
   ```

5. Confirm your configuration:

   ```
   show ip ospf accept
   ```

Example

Log on to the OSPF Router Configuration mode in ACLI:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config):router ospf
```

Create a policy to accept external routes from a specified advertising route:

```
Switch:1(config-ospf):accept adv-rtr 192.0.2.122
```

Enable an OSPF accept entry for a specified advertising route:

```
Switch:1(config-ospf):accept adv-rtr 192.0.2.122 enable
```

Exit to the Privileged EXEC mode:

```
Switch:1(config-ospf):exit
Switch:1(config):exit
```

Apply the OSPF accept policy change and confirm your configuration:

```
Switch:1#ip ospf apply accept adv-rtr 192.0.2.122
Switch:1#show ip ospf accept
================================================================================
                           Ospf Accept - GlobalRouter
================================================================================
ADV_RTR         MET_TYPE ENABLE POLICY
--------------------------------------------------------------------------------
192.0.2.122      -       FALSE
```

Variable definitions

Use the data in the following table to use the **`accept adv-rtr`** command.

Table 28: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the IP address.
enable	Enables an OSPF accept entry for a specified advertising router.

Variable	Value
	Use the no operator to disable an OSPF accept entry: `no accept adv-rtr <A.B.C.D> enable`
metric-type {type1|type2|}	Indicates the OSPF external type. This parameter describes which types of OSPF external routes match this entry. means match all external routes. *type1* means match external type 1 only. *type2* means match external type 2 only. Use the no operator to disable metric-type: `no ip ospf accept adv-rtr <A.B.C.D> metric-type`
route-policy <WORD>	Specifies the name of the route policy to use for filtering external routes advertised by the specified advertising router before accepting into the routing table.

Applying OSPF accept policy changes

Apply OSPF accept policy changes to allow the configuration changes in the policy to take effect in an OSPF Accept context (and to prevent the device from attempting to apply the changes one by one after each configuration change).

For more information on IS-IS accept policy filters for Avaya VENA Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

About this task

Important:

Changing OSPF Accept contexts is a process-oriented operation that can affect system performance and network accessibility while you perform the procedures. If you want to change the default preferences for an OSPF Accept or a prefix-list configuration (as opposed to the default preference), Avaya recommends that you do so before enabling the protocols.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Apply an OSPF accept policy change:

   ```
   ip ospf apply accept [vrf WORD<1-16>]
   ```

3. Display information about the configured OSPF entries:

   ```
   show ip ospf accept [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

Apply the OSPF accept policy and confirm the configuration:

```
Switch:1>enable
Switch:1#ip ospf apply accept
Switch:1#show ip ospf accept
================================================================================
                           Ospf Accept - GlobalRouter
================================================================================
ADV_RTR         MET_TYPE ENABLE POLICY
--------------------------------------------------------------------------------
192.0.2.122      -        TRUE
```

Variable definitions

Use the data in the following table to use the **ip ospf apply accept adv-rtr** command.

Table 29: Variable definitions

Variable	Value
adv-rtr	Commits entered changes. Issue this command after you modify a policy configuration that affects an OSPF accept policy.
vrf WORD<1–16>	Specifies the name of the VRF.

Use the data in the following table to use the **show ip ospf accept** command output.

Table 30: Variable definitions

Variable	Value
ADV_RTR	Indicates the router advancing the packets.
MET_TYPE	Indicates the metric type for the routes to import into OSPF routing protocol, which passed the matching criteria configured in this route policy. Options include: local, internal, external, externaltype1, and externaltype2.
ENABLE	Indicates if the policy is enabled.
POLICY	Indicates the type of policy.

Configuring inter-VRF redistribution policies

Configure redistribution entries to allow a protocol to announce routes of a certain source type, for example, static, RIP, or direct.

For more information on IS-IS redistribution, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Before you begin

- Ensure the routing protocols are globally enabled.

- You must configure the route policy, if required.
- Ensure the VRFs exist.
- You must create the route policy and prefix list under the source VRF context.
- You must log on to the VRF Router Configuration mode in ACLI.

Procedure

1. Create the redistribution instance:

   ```
   ip <rip|ospf|bgp> redistribute <ospf|bgp|static|direct|rip>
   ```

2. Apply a route policy if required:

   ```
   ip <rip|ospf|bgp> redistribute <ospf|bgp|static|direct|rip> route-
   policy <WORD 0-64> [vrf-src <WORD 1-16>]
   ```

3. Use the following variable definitions table to configure other parameters as required.
4. Enable the redistribution:

   ```
   ip <rip|ospf|bgp> redistribute <ospf|bgp|static|direct|rip> enable
   [vrf-src <WORD 1-16>]
   ```

5. Ensure that the configuration is correct:

   ```
   show ip <rip|ospf|bgp> redistribute [vrf WORD<1-16>] [vrfids
   WORD<0-512>]
   ```

6. Apply the redistribution:

   ```
   ip <rip|ospf|bgp> apply redistribute <ospf|bgp|static|direct|rip>
   [vrf WORD<1-16>] [vrf-src WORD<1-16>]
   ```

Example

```
Switch:1>enable
Switch:1#config terminal
```

Log on to the VRF Router Configuration mode:

```
Switch:1(config)#router vrf test
```

Create the redistribution instance:

```
Switch:1(router-vrf)#ip rip redistribute ospf
```

Enable the redistribution

```
Switch:1(router-vrf)#ip rip redistribute ospf enable
```

Ensure that the configuration is correct:

```
Switch:1(router-vrf)#show ip rip redistribute
```

Exit to Global Configuration mode:

```
Switch:1(router-vrf)#exit
```

Apply the redistribution:

```
Switch:1(config)#ip rip apply redistribute ospf
```

Variable definitions

Use the data in the following table to use the redistribution commands.

Table 31: Variable definitions

Variable	Value
<ospf|bgp|static|direct|rip>	Specifies the type of routes to redistribute—the protocol source.
vrf *WORD<1-16>*	Specifies the VRF instance.
vrfids *WORD<0-512>*	Specifies a list of VRF IDs.
vrf-src *WORD<1-16>*	Specifies the source VRF instance. This parameter is not required for redistribution within the same VRF.

Use the data in the following table to use the **`ip <bgp|ospf|rip> redistribute <ospf|bgp|static|direct|rip>`** command.

Variable	Value
apply [vrf-src WORD<1–16>]	Applies the redistribution configuration.
enable [vrf-src WORD<1–16>]	Enables the OSPF route redistribution instance.
metric <metric-value> [vrf-src WORD<1–16>]	Configures the metric to apply to redistributed routes.
metric-type <type1|type2> [vrf-src WORD<1–16>]	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
route-policy <policy-name> [vrf-src WORD<1–16>]	Configures the route policy to apply to redistributed routes.
subnets <allow|suppress> [vrf-src WORD<1–16>]	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.

Chapter 9: IP policy configuration using Enterprise Device Manager

You can form a unified database of route policies that the protocols (RIP, OSPF or Border Gateway Protocol [BGP]) can use for any type of filtering task.

For information about configuring a prefix list, community list, or AS path list, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

A name or an ID identifies a policy. Under a policy you can have several sequence numbers, each of which is equal to one policy in the old convention. If a field in a policy is not configured, it appears as 0 or any when it appears in Enterprise Device Manager (EDM). This means that the field is ignored in the match criteria. You can use the clear option to remove existing configurations for any field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a set-preference field set, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce and redistribute purposes.

You can apply only one policy for one purpose (for example, RIP Announce on a given RIP interface). In that example, all sequence numbers under the given policy are applicable for that filter. A sequence number also acts as an implicit preference: a lower sequence number is preferred.

Configuring a prefix list

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

Before you begin

- Change the VRF instance as required to configure a prefix list on a specific VRF instance.

Procedure

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

3. Click the **Prefix List** tab.
4. Click **Insert**.
5. In the **Id** box, type an ID for the prefix list.
6. In the **Prefix** box, type an IP address for the route.
7. In the **PrefixMaskLength** box, type the length of the prefix mask.
8. Configure the remaining parameters as required.
9. Click **Insert**.

Prefix List field descriptions

Use the data in the following table to use the **Prefix List** tab.

Name	Description
Id	Configures the list identifier.
Prefix	Configures the IP address of the route.
PrefixMaskLen	Configures the specified length of the prefix mask. You must enter the full 32-bit mask to exact a full match of a specific IP address, for example, if you create a policy to match on the next hop.
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name length can use from 1 to 64 characters.
MaskLenFrom	Configures the lower bound of the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Configures the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Configuring a route policy

Configure a route policy so that all protocols use them for In, Out, and Redistribute purposes.

For more information on IS-IS accept policy filters for Avaya VENA Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

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

3. Click the **Route Policy** tab.
4. Click **Insert**.
5. Enter the appropriate information for your configuration in the Insert Route Policy dialog box.
6. Click **Insert**.

Route Policy field descriptions

Use the data in the following table to use the **Route Policy** tab.

Name	Description
Id	Specifies the ID of an entry in the Prefix list table.
SequenceNumber	Specifies a policy within a route policy group.
Name	Specifies the name of the policy. This command changes the name field for all sequence numbers under the given policy.
Enable	Indicates whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored. The default is disabled.
Mode	Specifies the action to take when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route). The default is permit.
MatchProtocol	Selects the appropriate protocol. If configured, matches the protocol through which the route is learned. This field is used only for RIP Announce purposes. The default is to enable all match protocols.
MatchNetwork	Specifies if the system matches the destination network against the contents of the specified prefix list.
MatchIpRouteSource	Specifies if the system matches the next-hop 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. Click the ellipsis button and choose from the list in the Match Route Source dialog box. You can select up to four entries. To clear an entry, use the ALT key. You can also change this field in the Route Policy tab of the Policy dialog box.
MatchNextHop	Specifies if the system matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes. Click the ellipsis button and choose from the list in the Match Next Hop dialog box. You can select up to four entries. To clear an entry, use the ALT key.

Name	Description
MatchInterface	Specifies if the system 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. Click the ellipsis button and choose from the list in the Match Interface dialog box. You can select up to four entries. To clear an entry, use the ALT key.
MatchRouteType	Configures a specific route type to match (applies only to OSPF routes). Externaltype1 and Externaltype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes. The default is any.
MatchMetric	Specifies if the system matches the metric of the incoming advertisement or existing route against the specified value (1 to 65535). If 0, this field is ignored. The default is 0.
MatchAsPath	Configures if the system matches the BGP autonomus system path. Applicable to BGP only. This overrides the BGP neighbor filter list information.
MatchCommunity	Filters incoming and outgoing updates based on a Community List. Applicable to BGP only. The default is disable.
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.
MatchTag	Specifies a list of tags used during the match criteria process. Applicable to BGP only. It contains one or more tag values.
MatchVrf	Identifies the source VRFs that leaks routes to the local VRF (applies only to RIP routes).
NssaPbit	Configures or resets the P bit in specified type 7 link state advertisements (LSA). By default, the Pbit is always configured in case the user configures the Pbit to a disable state for a particular route policy other than all type 7. LSAs associated with that route policy have the Pbit cleared. With this intact, not so stubby area (NSSA) area border router (ABR) does not perform translation of these LSAs to type 5. The default is enable.
SetRoutePreference	Configures a value from 0 to 255. The default value is 0. If the default is configured, the global preference value is used. When configured to a value greater than zero, specifies the route preference value assigned to the routes that matches the policy. This feature applies to accept policies only.
SetMetricTypeInternal	Identifies the MED value for routes advertised to BGP numbers to the Interior Gateway Protocol (IGP) metric value. The default is 0.

Name	Description
SetMetric	Configures the system to use 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 (applies to IS-IS routes also). The default is 0.
SetMetricType	Configures the metric type for the routes to announce into the OSPF routing protocol that matches this policy. Applicable to OSPF protocol only. The default is type 2. This field is applicable only for OSPF announce policies. The default is type2.
SetNextHop	Configures the IP address of the next-hop router. Applicable to BGP only. The default is 0.0.0.0.
SetInjectNetList	Configures the destination network of the route that matches this policy with the contents of the specified prefix list. Click the ellipsis button and choose from the list in the Set Inject NetList dialog box.
SetMask	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.
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 The default is prepend.
SetAutomaticTag	Enables the automatic tag feature. Applicable to BGP protocol only. The default is disable.
SetCommunityNumber	Configures the community number for BGP advertisements. This value can be a number (1 to 42949672000) or no-export or no-advertise.
SetCommunityMode	Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged. • 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.

Name	Description
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.
SetOriginEgpAs	Indicates the remote autonomous system number for the BGP protocol. The default is 0.
SetWeight	Configures the weight value for the routing table for the BGP protocol. This field must be used with the match as-path condition. For BGP, this value overrides the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. The default is 0.
SetTag	Configures the list of tags used during the match criteria process for the BGP protocol. The default is 0.
Ipv6SetNextHop	Specifies the address of the IPv6 next hop router.

Applying a route policy

Apply route policies to define route behavior.

For more information on IS-IS accept policy filters for Avaya VENA Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

About this task

Important:

Changing route policies or prefix lists that affect OSPF accept or redistribute is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. Therefore, Avaya recommends that if you want to change a prefix list or a routing protocol, you configure all route policies and prefix lists before enabling the protocols.

Procedure

1. In the navigation tree, expand the following folders:**Configuration** > **IP**.
2. Click **Policy**
3. Click the **Applying Policy** tab.
4. Select the type of policy to apply.
5. Click **Apply**.

Applying Policy field descriptions

Use the data in the following table to use the **Applying Policy** tab.

Name	Description
RoutePolicyApply	Specifies that configuration changes in the policy take effect in an OSPF route policy context. This prevents the system from attempting to apply the changes one by one after each configuration change. The default is enabled.
RedistributeApply	Specifies that configuration changes in the policy take effectfor an OSPF Redistribute context. This prevents the system from attemptingto apply the changes one-by-one after each configuration change. The default is enabled.
OspfInFilterApply	Specifies that configuration changes in a route policy or a prefix list take effect in an OSPF Accept context. This prevents the system from attempting to apply the changes one by one after each configuration change. The default is enabled.

Viewing IP routes

View IP routes learned on the device.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Routes** tab to view IP routes learned on the device.
4. If you want to limit the routes displayed, click **Filter** to show a smaller subset of the learned routes.
5. In the Filter dialog box, select an option, or options, and enter information to limit the routes to display in the Routes table.
6. Click **Filter** and the Routes table displays only the routes that match the options and information that you enter.

Routes field descriptions

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

Name	Description
Dest	Specifies the destination IP network of this route. An entry with a value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to multiple entries depends on the table

Name	Description
	access mechanisms defined by the network management protocol in use.
Mask	Indicates the network mask to logically add with the destination address before comparison to the destination IP network.
NextHop	Specifies the IP address of the next hop of this route.
NextHopId	Displays the MAC address or hostname of the next hop.
HopOrMetric	Displays the primary routing metric for this route. The semantics of this metric are specific to different routing protocols.
Interface	Specifies the router interface for this route. • Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation. • Brouter interfaces are identified by the slot and port number of the brouter port.
Proto	Specifies the routing mechanism through which this route was learned: • other—none of the following • local—nonprotocol information, for example, manually configured entries • static • ICMP • EGP • GGP • Hello • RIP • IS-IS • ES-IS • Cisco IGRP • bbnSpfIgp • OSPF • BGP • Inter-VRF Redistributed Route
Age	Displays the number of seconds since this route was last updated or otherwise determined to be correct.
PathType	Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths. • iA indicates Indirect Alternative route without an ECMP path • iAE indicates Indirect Alternative ECMP path • iB indicates Indirect Best route without ECMP path

Name	Description
	• iBE indicates Indirect Best ECMP path • dB indicates Direct Best route • iAN indicates Indirect Alternative route not in hardware • iAEN indicates Indirect Alternative ECMP route not in hardware • iBN indicates Indirect Best route not in hardware • iBEN indicates Indirect Best ECMP route not in hardware • dBN indicates Direct Best route not in hardware • iAU indicates Indirect Alternative Route Unresolved • iAEU indicates Indirect Alternative ECMP Unresolved • iBU indicates Indirect Best Route Unresolved • iBEU indicates Indirect Best ECMP Unresolved • dBU indicates Direct Best Route Unresolved • iBF indicates Indirect Best route replaced by FTN • iBEF indicates Indirect Best ECMP route replaced by FTN • iBV indicates Indirect best IPVPN route • iBEV indicates Indirect best ECMP IP VPN route • iBVN indicates Indirect best IP VPN route not in hardware • iBEVN indicates Indirect best ECMP IP VPN route not in hardware
Pref	Displays the preference.
Layer3VirtualInterfaceType	Identifies the type for the value in the **Layer3VirtualInterface** field. The values include: • none – Specifies the type is not applicable for the route. • spb – Specifies that the routes are learned through ISIS and SPBM.
Layer3VirtualInterface	Specifies the layer 3 virtual interface. The values include: • 0 – Specifies the Global Router. • -1 – Specifies the route is not applicable.

Configuring an OSPF accept policy

Perform the following procedure to create or configure an OSPF accept policy.

For more information on IS-IS accept policy filters for Avaya VENA Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click the **OSPF Accept** tab.
4. Click **Insert**.
5. Configure the parameters as required.
6. Click **Insert**.

OSPF Accept field descriptions

Use the data in the following table to use the **OSPF Accept** tab.

Name	Description
AdvertisingRtr	Specifies the routing ID of the advertising router.
Enable	Enables or disables the advertising router. You can also enable or disable advertising in the OSPF Accept tab of the Policy dialog box by clicking in the field and selecting enable or disable from the menu. The default is disable.
MetricType	Specifies the OSPF external type. This parameter describes which types of OSPF ASE routes match this entry. • Any means match either ASE type 1 or 2 • Type1 means match any external type 1 • Type2 means match any external type 2 You can also select your entry in the OSPF Accept tab of the Policy dialog box by clicking in the field and selecting any, type1, or type2 from the menu. The default is any.
PolicyName	Specifies the name of the OSPF in filter policy. Click the ellipsis button and choose from the list in the Policy Name dialog box. To clear an entry, use the ALT key.

Configuring inbound/outbound filtering policies on a RIP interface

About this task

Configure inbound filtering on a RIP interface to determine whether to learn a route on a specified interface and to specify the parameters of the route when it is added to the routing table. Configure

outbound filtering on a RIP interface to determine whether to advertise a route from the routing table on a specified interface and to specify the parameters of the advertisement.

The port on which the multimedia filter is enabled becomes a DIFFSERV access port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click the **RIP In/Out Policy** tab.
4. In the desired row, double-click the **InPolicy** or **OutPolicy** column.
5. Select a preconfigured In/Out policy and click **OK**.

RIP In/Out Policy field descriptions

Use the data in the following table to use the **RIP In/Out Policy** tab.

Name	Description
Address	Specifies the IP address of the RIP interface.
Interface	Specifies the internal index of the RIP interface.
InPolicy	Specifies the policy name used for inbound filtering on this RIP interface. This policy determines whether to learn a route on this interface and specifies the parameters of the route when it is added to the routing table.
OutPolicy	Specifies the policy name used for outbound filtering on this RIP interface. This policy determines whether to advertise a route from the routing table on this interface and specifies the parameters of the advertisement.

Deleting inbound/outbound filtering policies on a RIP interface

About this task

Delete a RIP In/Out policy when you no longer want to learn a route on a specified interface or advertise a route from the routing table on a specified interface.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click the **RIP In/Out Policy** tab.
4. In the desired row, double-click the **InPolicy** or **OutPolicy** column for the policy you want to delete.

5. In the **InPolicy** or **OutPolicy** dialog box, press `CTRL` and then, click the policy you want to delete.
6. Click **OK**.

   The policy is deleted and you are returned to the RIP In/Out Policy tab.
7. Click **Apply**.

Chapter 10: IP routing configuration using ACLI

Configure the IP router interface so that you can configure and use routing protocols and features on the interface. This section contains instructions for both the Global Router and Virtual Router Forwarding (VRF) instances.

Enabling routing globally or on a VRF instance

Use IP forwarding (routing) on a global level so that the device supports routing. You can use the IP address of an interface for IP-based network management.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Activate IP forwarding:

   ```
   ip routing
   ```

3. View the forwarding configuration:

   ```
   show ip routing [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
CB-SWA:1(config)#show ip routing

================================================================
                              IP - GlobalRouter
================================================================


 IP Forwarding is enabled
 IP ECMP feature is disabled
 Maximum ECMP paths number is 1
 ECMP 1 pathlist :
 ECMP 2 pathlist :
```

```
ECMP 3 pathlist :
ECMP 4 pathlist :
ECMP 5 pathlist :
ECMP 6 pathlist :
ECMP 7 pathlist :
ECMP 8 pathlist :
IP Alternative Route feature is enabled
IP More Specific Non Local Route feature is disabled
IP ICMP Unreachable Message is disabled
IP Supernetting is disabled
IP Icmp-redirect-msg is disabled
IP Default TTL is 255 seconds

IP ARP life time is 360 minutes
```

Variable definitions

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

Table 32: Variable definitions

Variable	Value
vrf *WORD<1-16>*	Specifies a VRF instance by VRF name.
vrfids *WORD<0-512>*	Specifies a VRF instance by VRF number.

Enabling routing on a port

Before you begin

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

About this task

You can enable or disable routing capabilities on specified device ports. The specified port can be part of a routed VLAN, while routing is disabled only on that port. The default setting for routing is enable.

Procedure

Enable routing:

```
routing enable
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 3/1
VSP-9012:1(config-if)#routing enable
```

Deleting a dynamically learned route

Delete a dynamically learned route from the routing table if you do not want Virtual Services Platform 9000 to use the route. Exercise caution when you delete entries from the routing table.

About this task

The NH VRF/ISID column in the **`show ip route [vrf WORD<1-16>]`** command displays the I-SID in the following examples:

- Only for inter-Virtual Services Network (VSN) routes leaked using IS-IS accept policies.
- Only if the I-SID for which the routes are leaked does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays.

If the I-SID is 0, which represents the GlobalRouter, the column displays as GlobalRouter. The existing IS-IS routes in Shortest Path Bridging (SPB) Layer 3 VSN continue to display as the VRF name of the IP VSN.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. View IP route information:

   ```
   show ip route [<A.B.C.D>] [-s default|-s <A.B.C.D/X>] [alternative]
   [count-summary] [preference] [vrf WORD<1-16>] [vrfids WORD<0-512>]
   [static]
   ```

3. Delete the dynamically learned route:

   ```
   no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> dynamic
   ```

Example

Display IP route information. Delete a dynamically learned route.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#show ip route

================================================================================
                          IP Route - GlobalRouter
================================================================================
                                            NH                INTER
DST             MASK            NEXT        VRF/ISID     COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
198.51.100.1    255.255.255.255 192.0.2.65  GlobalRouter     1   100  OSPF 0   IB   125
198.51.100.5    255.255.255.255 192.0.2.5   -                1   0    LOC  0   DB   0
198.51.100.13   255.255.255.255 VSP13       GlobalRouter     10  1000 ISIS 0   IBS  7
198.51.100.200  255.255.255.255 VSP200      GlobalRouter     10  1000 ISIS 0   IBS  7

VSP-9012:1(config)#no ip route 198.51.100.1 255.255.255.0 192.0.2.65 dynamic
```


Comments? infodev@avaya.com

Variable definitions

Use the data in the following table to use the **show ip route** commands.

Table 33: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the IP address of the route to the network.
alternative	Displays the alternative routes.
count-summary	Displays a summary of the number of routes learned from each routing protocol for each VRF.
preference	Displays the route preference.
-s <A.B.C.D/X>	Indicates the IP address and subnet mask for which to display routes.
-s default	Indicates the default subnet.
static	Displays the static route information.
vrf *WORD<1-16>*	Displays the route for a particular VRF.
vrfids *WORD<0-512>*	Displays the route for a particular VRF number.

Use the data in the following table to use the **no ip route** command.

Table 34: Variable definitions

Variable	Value
<A.B.C.D> <A.B.C.D> <A.B.C.D>	Specifies the IP address, the subnet mask, and the next-hop IP address, respectively.
dynamic	Specifies that a dynamic route is to be deleted.
enable	Disables the route.
local-next-hop enable	Disables the local-next-hop option.
preference	Deletes the value of the route preference.
next-hop-vrf WORD<1-16>	Specifies the name of the next-hop VRF router.

Configuring IP route preferences

Configure IP route preferences to override default route preferences and give preference to routes learned for a specific protocol. You must disable ECMP before you configure route preferences.

Before you begin

- Ensure that ECMP is disabled.

Important:

Changing route preferences can affect system performance and network accessibility while you perform the procedure. Avaya recommends that you change a prefix list or a routing protocol before you activate the protocols.

About this task

To configure route preferences for a VRF, access VRF Router Configuration mode, rather than Global Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the route preference:

   ```
   ip route preference protocol <static|ospf-intra|ospf-inter|ebgp|
   ibgp|rip|ospf-extern1|ospf-extern2|staticv6|ospfv3-intra|ospfv3-
   inter|ospfv3-extern1|ospfv3-extern2|spbm-level1> <0-255>
   ```

3. Confirm that the configuration is correct:

   ```
   show ip route preference [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Configure the route preference to SPBM Level 1:

```
VSP-9012:1(config)#ip route preference protocol spbm-level1 7
```

Confirm the configuration is correct:

```
VSP-9012:1(config)# show ip route preference vrf test
```

```
================================================================================
                          IP Route Preference - VRF test
================================================================================
PROTOCOL        DEFAULT    CONFIG
--------------------------------------------------------------------------------
LOCAL           0          0
STATIC          5          5
OSPF_INTRA      20         20
OSPF_INTER      25         25
EBGP            45         45
RIP             100        100
OSPF_E1         120        120
OSPF_E2         125        125
IBGP            175        175
STATICv6        5          5
OSPFv3_INTRA    20         20
OSPFv3_INTER    25         25
OSPFv3_E1       120        120
OSPFv3_E2       125        125
SPBM_L1         7          7
```

Variable definitions

Use the data in the following table to use the **ip route preference** and the **show ip route preference** commands.

Table 35: Variable definitions

Variable	Value
protocol <static|ospf-intra|ospf-inter|ebgp|ibgp|rip|ospf-extern1|ospf-extern2|staticv6|ospfv3-intra|ospfv3-inter|ospfv3-extern1|ospfv3-extern2|spbm-level1> <0-255>	Configures the preference value for the specified protocol. If two protocols have the same configured value, the default value is used. • The protocol must be one of the following: static, ospf-intra, ospf-inter, ebgp, ibgp, rip, ospf-extern1, ospf-extern2, staticv6, ospfv3-intra, ospfv3-inter, ospfv3-extern1, ospfv3-extern2 or spbm-level1. • *<0-255>* configures the priority. 0 is reserved for local routes. The default is 7.
vrf WORD<1-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-512>	Specifies a VRF instance by VRF number.

Flushing routing tables by VLAN or port

Before you begin

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

About this task

For administrative and troubleshooting purposes, flush the routing tables.

To flush tables on a VRF instance for a port or VLAN, ensure that the VRF is associated with the port or VLAN.

Procedure

Flush the routing tables:

```
action flushIp
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 3/15
VSP-9012:1(config-if)#action flushIp
```

Configuring an IP address for the management port

Before you begin

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

About this task

Configure the IP address for the management port so that you can remotely access the device using the management port. The management port runs on a dedicated VRF and Avaya recommends that you redirect all commands that are run on the management port to its VRF.

The configured IP subnet has to be globally unique because the management protocols can go through in-band or out-of-band ports.

Procedure

1. Configure the IP address and mask for the management port:

   ```
   ip address <A.B.C.D> <A.B.C.D>
   ```

2. Show the complete network management information:

   ```
   show interface mgmtEthernet
   ```

3. Show the management IP interface information:

   ```
   show ip interface vrf mgmtrouter
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface mgmtethernet 2/1
```

Configure the IP address for the management port:

```
VSP-9012:1(config-if)#ip address 47.17.10.31 255.255.255.0
```

Variable definitions

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

Table 36: Variable definitions

Variable	Value
<A.B.C.D> <A.B.C.D>	Specifies the IP address followed by the subnet mask.

Configuring a virtual IP address for the management port

Before you begin

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

About this task

Configure a virtual IP address for the management port so that you have an alternative way to remotely connect to it. After you change the boot configuration, you must save the changes to both the Master and the Standby CPUs.

The virtual IP address for the management port is a floating address that is always owned by the current primary CPU. After a switchover, the IP address is automatically assigned to the new primary CPU, which allows you to use a single IP address to access the management interface regardless of which CPU is the primary.

Procedure

1. Configure a virtual IP address for the management port:

   ```
   sys mgmt-virtual-ip <A.B.C.D/X>
   ```

2. Save the configuration to the master CPU:

   ```
   save config
   ```

3. If the boot config flag savetostandby is false, save the configuration to the standby CPU:

   ```
   save config standby
   ```

4. Use Telnet to connect to the standby CPU:

   ```
   peer telnet
   ```

5. Reset the standby CPU:

   ```
   reset
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#sys mgmt-virtual-ip 0.0.0.0/0.0.0.0
VSP-9012:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **`sys mgmt-virtual-ip`** command.

Table 37: Variable definitions

Variable	Value
<A.B.C.D/X>	Specifies the IP address and network mask (0–32) that you assign to the network management port. The default form of this command is `default sys mgmt-virtual-ip`. The no form of this command is `no sys mgmt-virtual-ip`.

Assigning an IP address to a port

Assign an IP address to a port so that it supports routing operations.

Use a brouter port to route IP packets and to bridge all nonroutable traffic. 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 nonroutable traffic and still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

Before you begin

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

About this task

If an IP interface is configured without specifying the VRF instance, it maps to VRF 0 by default.

Use the `vrf` parameter to associate the port or VLAN with a VRF instance.

Procedure

1. Assign an IP address to the port:

   ```
   brouter port {slot/port} vlan <2-4084> subnet <A.B.C.D/X> [mac-offset <0-65535>]
   ```

2. If required, associate the port with a VRF:

   ```
   vrf WORD<1-16>
   ```

3. Confirm that the configuration is correct:

   ```
   show brouter [<1-4084>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 3/11
```

Assign an IP address to the port

```
VSP-9012:1(config)#brouter port 3/11 vlan 2202 subnet
47.17.10.31/255.255.255.0
```

Variable definitions

Use the data in the following table to use the **brouter port** command.

Table 38: Variable definitions

Variable	Value
mac-offset <0-65535>	Specifies a number by which to offset the MAC address of the brouter port from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated.
slot/port	Indicates the slot and port number of the port you are configuring.
subnet <A.B.C.D/X>	Specifies the IP address and subnet mask (0–32).
<2-4084>	Specifies the VLAN ID that is used if the port is tagged (802.1q encapsulation). The VLAN ID is unique to Virtual Services Platform 9000 and is not used if the port is untagged.

Use the data in the following table to use the **show brouter** command.

Table 39: Variable definitions

Variable	Value
<1-4084>	Specifies the VLAN ID that is used if the port is tagged (802.1q encapsulation). The VLAN ID is unique to Virtual Services Platform 9000 and is not used if the port is untagged.

Assigning an IP address to a VLAN

Assign an IP address to a VLAN so that it supports routing operations.

Before you begin

- Activate IP forwarding globally.

About this task

If an IP interface is configured without specifying the VRF instance, it maps to VRF 0 by default.

Use the vrf parameter to associate the VLAN with a VRF instance.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4084>
```

2. Assign an IP address:

```
ip address {A.B.C.D} {A.B.C.D} [<0-65535>]
```

3. If required, associate the VLAN with a VRF:

```
vrf WORD<1-16>
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 2
VSP-9012:1(config-if)#ip address 47.17.10.32 255.255.255.0
```

Variable definitions

Use the data in the following table to complete the **ip address** commands.

Table 40: Variable definitions

Variable	Value
<A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask, respectively.
<0-65535>	mac-offset specifies a number by which to offset the MAC address of the brouter port or VLAN from the chassis MAC address. This ensures that each IP address has a different MAC address. The range is 0–65535.

Use the data in the following table to use the **vrf** command.

Table 41: Variable definitions

Variable	Value
WORD<1-16>	Specifies the VRF of the VLAN.

Viewing IP addresses for all router interfaces

Before you begin

- You must log on to the Privileged EXEC mode in ACLI.

About this task

Perform the following procedure to display information about all IP interfaces configured on the device.

Procedure

Show the IP interfaces and addresses on the device:

```
show ip interface
```

Example

```
VSP-9012:1>enable
VSP-9012:1#show ip interface

================================================================================
                          IP Interface - GlobalRouter
================================================================================
INTERFACE    IP              NET             BCASTADDR  REASM    VLAN  BROUTER
             ADDRESS         MASK            FORMAT     MAXSIZE  ID    PORT
--------------------------------------------------------------------------------
Port3/11     20.20.20.20     255.255.255.0   ones       1500     2202  true
Vlan2        2.2.2.31        255.0.0.0       ones       1500     2     false
Vlan3        30.30.30.31     255.255.255.0   ones       1500     3     false
Vlan4        40.1.1.31       255.255.255.0   ones       1500     4     false
Vlan5        45.0.0.1        255.0.0.0       ones       1500     5     false
Vlan6        50.0.0.1        255.0.0.0       ones       1500     6     false
Vlan10       10.10.10.2      255.0.0.0       ones       1500     10    false
Vlan100      100.1.1.1       255.255.255.0   ones       1500     100   false
Vlan200      1.1.1.1         255.255.255.0   ones       1500     200   false
Vlan400      4.1.1.31        255.255.255.0   ones       1500     400   false


All 10 out of 10 Total Num of IP interfaces displayed
```

Variable definitions

Use the data in the following table to **`show ip interface`** command.

Table 42: Variable definitions

Variable	Value
gigabitethernet*[{slot/port[-slot/port][,...]}]*	Displays IP interface information for Gigabit Ethernet ports.
vrf*WORD<1–16>*	Displays interface information for a particular VRF.
vrfids*WORD<0–512>*	Displays interface information for particular VRF IDs.

Configuring IP routing globally or for a VRF

Configure the IP routing protocol stack to specify which routing features the device can use. You can configure global parameters before or after you configure the routing protocols.

About this task

To configure IP routing globally for a VRF instance, use VRF Router Configuration mode rather than Global Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the default TTL for all routing protocols to use:

   ```
   ip ttl <1-255>
   ```

   This value is placed into routed packets that have no TTL specified.

3. Activate ECMP:

   ```
   ip ecmp
   ```

4. Activate the alternative route feature globally:

   ```
   ip alternative-route
   ```

5. Configure a prefix-list for target destination:

   ```
   ip prefix-list WORD<1-64> <A.B.C.D/X>
   ```

6. Set ECMP prefix-list to specify routes with needed number of paths:

   ```
   ip ecmp pathlist-<1-8> WORD<1-64>
   ```

7. Access privileged EXEC mode:

   ```
   end
   ```

8. Apply changes to all ECMP path-list apply configurations:

   ```
   ip ecmp pathlist-apply
   ```

9. Configure the remaining global parameters as required.

10. Display the list of routes with number of ECMP paths.

    ```
    show ip ecmp max-path [vrf WORD<1-16>][vrfids WORD<0-512>]
    ```

Example

Enable ECMP. Configure prefix-list LIST1 for the target destination of 47.17.121.50/255.255.255.0. Configure the ECMP prefix-list to specify routes with the needed number of paths. Apply changes to all ECMP path-list apply configurations. Display the prefix list of routes for a particular VRF ID.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#ip ecmp
VSP-9012:1(config)#ip prefix-list LIST1 47.17.121.50/255.255.255.0
VSP-9012:1(config)#ip ecmp pathlist-1 LIST1
VSP-9012:1(config)#end
VSP-9012:1(config)#ip ecmp pathlist-apply
```

```
VSP-9012:1(config)#show ip ecmp max-path vrfids 12

=============================================================================
                        ecmp-max-path - VRF "virtualrandf12"
=============================================================================


                ecmp-max-path : 1
```

Variable definitions

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

Table 43: Variable definitions

Variable	Value
alternative-route	Enables or disables the alternative route feature. The default value is enabled. If the alternative-route parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are readded. The default form of this command is `default ip alternative-route`. The no form of this command is `no ip alternative-route`.
max-routes-trap enable	Enables the device to send a trap after the maximum number of routes is exceeded. The no form of this command is `no max-routes-trap enable`. The default form of this command is `default max-routes-trap enable`.
more-specific-non-local-route	Enables the more-specific-non-local-route feature. If enabled, the device can enter a more-specific nonlocal route into the routing table. The default is disabled. The default form of this command is `default ip more-specific-non-local-route`. The no form of this command is `no ip more-specific-non-local-route`.
routing	Enables routing. The no form of this command is `no ip routing`.
supernet	Enables or disables supernetting. If you globally enable supernetting, the device can learn routes with a route mask of less then eight bits. Routes with a mask length less than eight bits cannot have ECMP paths, even if the ECMP feature is globally enabled. The default is disabled.

Variable	Value
	The default form of this command is `default ip supernet`. The no form of this command is `no ip supernet`.
ttl *<1-255>*	Configures the default time-to-live (TTL) value for a routed packet. The TTL is the maximum number of seconds before a packet is discarded. The default value of 255 is used whenever a time is not supplied in the datagram header. The default form of this command is `default ip ttl`.

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

Table 44: Variable definitions

Variable	Value
pathlist-1 *WORD<0-64>*	Configures one equal-cost path to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-1`.
pathlist-2 WOR<0-64> *WORD<0-64>*	Configures up to two equal-cost paths to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-2`.
pathlist-3 *WORD<0-64>*	Configures up to three equal-cost paths to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-3`.
pathlist-4 *WORD<0-64>*	Configures up to four equal-cost paths to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-4`.
pathlist-5 *WORD<0-64>*	Configures up to five equal-cost paths to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-5`.
pathlist-6 *WORD<0-64>*	Configures up to six equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.

Variable	Value
	To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-6`.
pathlist-7 *WORD<0-64>*	Configures up to seven equal-cost paths to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-7`.
pathlist-8 *WORD<0-64>*	Configures up to eight equal-cost paths to the same destination prefix. To remove the policy, enter a blank string. To configure this parameter, you must globally activate ECMP, with the command `ip ecmp`. The no form of this command is `no ip ecmp pathlist-8`.
max-path *<1-8>*	Configures the maximum number of ECMP paths. The range for this number 1–8. The default form of this command is `default ip ecmp max-path`.

Use the data in the following table to use the **`ip icmp`** commands.

Table 45: Variable definitions

Variable	Value
redirect	Enables the device to send ICMP destination redirect messages. The default form of this command is `default ip icmp redirect`.
unreachable	Enables the device to send ICMP unreachable messages. When enabled, this variable generates Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this router. These messages help determine if the device is reachable over the network. The default is disabled. The default form of this command is `default ip icmp unreachable`.

Use the data in the following table to use the **`show ip icmp`** commands.

Variable	Value
max-path	Displays the maximum number of Equal Cost Multipath (ECMP) paths.
vrf *WORD<1–16>*	Displays the prefix list of routes for a particular VRF. WORD<1–16> specifies the VRF name.
vrfids *WORD<0–512>*	Displays the prefix list of routes for a particular VRF ID. WORD<0–512> specifies the VRF ID.

Configuring static routes

Before you begin

- You must log on to the Global Configuration mode or the VRF Router Configuration mode in ACLI.
- Ensure no black hole static route exists.

About this task

Configure a static route when you want to manually create a route to a destination IP address.

If a black hole route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

Procedure

1. Create an IP static route:

   ```
   ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> weight <1-65535>
   ```

2. Enable an IP static route:

   ```
   ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> enable
   ```

3. Use the following variable definitions table to configure other static route parameters as required.
4. View existing IP static routes for the device, or for a specific network or subnet:

   ```
   show ip route static
   ```

5. Delete a static route:

   ```
   no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D>
   ```

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Create an IP static route:

```
VSP-9012:1(config)#ip route 42.17.0.0 255.255.0.0 42.17.156.126 weight
200
```

Enable a static route:

```
VSP-9012:1(config)#ip route 42.17.0.0 255.255.0.0 42.17.156.126 enable
```

View existing IP static routes for the device, or for a specific network or subnet:

```
VSP-9012:1(config)#show ip route static
```

Variable definitions

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

Table 46: Variable definitions

Variable	Value
<A.B.C.D> <A.B.C.D> <A.B.C.D>	The first and second <A.B.C.D> specify the IP address and mask for the route destination. The third <A.B.C.D> specifies the IP address of the next-hop router (the next router at which packets must arrive on this route). When you create a black hole static route, configure this parameter to 255.255.255.255 as the IP address of the router through which the specified route is accessible.
enable	Adds a static or default route to the router or VRF. The no form of this command is `no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> enable`. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> enable`.
local-next-hop enable	Enables the local next hop for this static route. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> local-next-hop enable`. The no form of this command is `no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> local-next-hop enable`.
next-hop-vrf *WORD<1-16>*	Specifies the next-hop VRF instance by name. After you configure the next-hop-vrf parameter, the static route is created in the local VRF, and the next-hop route is resolved in the next-hop VRF instance (next-hop-vrf). The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> next-hop-vrf WORD<1-16>`. The no form of this command is `no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> next-hop-vrf WORD<1-16>`.
weight *<1-65535>*	Specifies the static route cost. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> weight`.
preference *<1-255>*	Specifies the route preference. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> preference`.

Use the data in the following table to use the **show ip route static** command.

Table 47: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the route by IP address.
-s { *<A.B.C.D>* *<A.B.C.D>* | default}	Specifies the route by IP address and subnet mask.
vrf *WORD<1-16>*	Specifies a VRF by name.
vrfids *WORD<0-512>*	Specifies a range of VRF IDs.

Configuring a black hole static route

Configure a black hole static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.

If a black hole route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create a black hole static route:

   ```
   ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 weight <1-65535>
   ```

3. Enable a black hole static route:

   ```
   ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 enable [next-hop-vrf
   WORD<1-16>]
   ```

4. Configure other black hole static route parameters as required.

   When you specify a route preference, appropriately configure the preference so that when the black-hole route is used, it is elected as the best route.

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Create a black hole static route:

```
VSP-9012:1(config)#ip route 42.17.0.0 255.255.0.0 255.255.255.255 weight
200
```

Enable a black hole static route:

```
VSP-9012:1(config)#ip route 42.17.0.0 255.255.0.0 255.255.255.255 enable
```

Variable definitions

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

Table 48: Variable definitions

Variable	Value
<A.B.C.D>	The first and second *<A.B.C.D>* specify the IP address and mask for the route destination. 255.255.255.255 is the destination of the black hole route.
enable	Adds a static or default route to the router or VRF. The no form of this command is `no ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 enable`.
local-next-hop enable	Enables the local next hop for this static route. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> local-next-hop enable`. The no form of this command is `no ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 local-next-hop enable`.
next-hop-vrf *WORD<1-16>*	Specifies the next-hop VRF instance by name. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 next-hop-vrf WORD<1-16>`. The no form of this command is `no ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 next-hop-vrf WORD<1-16>`.
weight *<1-65535>*	Specifies the static route cost. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 weight`.
preference *<1-255>*	Specifies the route preference. The default form of this command is `default ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 preference`.

Configuring a default static route

Use the default route to specify a route to all networks for which there are no explicit routes in the forwarding information base or the routing table. This route has a prefix length of zero (RFC 1812). You can configure Virtual Services Platform 9000 systems with a static default route, or they can learn it through a dynamic routing protocol.

To create a default static route, you configure the destination address and subnet mask to 0.0.0.0.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create a default static route:

   ```
   ip route 0.0.0.0 0.0.0.0 <A.B.C.D> weight <1-65535>
   ```

3. Enable a default static route:

   ```
   ip route 0.0.0.0 0.0.0.0 <A.B.C.D> enable [next-hop-vrf WORD<1-16>]
   ```

4. Configure other default static route parameters as required.

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Create a default static route:

```
VSP-9012:1(config)#ip route 0.0.0.0 0.0.0.0 42.17.159.128 weight 100
```

Enable a default static route:

```
VSP-9012:1(config)#ip route 0.0.0.0 0.0.0.0 42.17.159.128 enable
```

Variable definitions

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

Table 49: Variable definitions

Variable	Value
<A.B.C.D>	<A.B.C.D> specifies the IP address of the next-hop router (the next router at which packets must arrive on this route).
enable	Adds a static or default route to the router or VRF. The no form of this command is `no ip route 0.0.0.0 0.0.0.0 <A.B.C.D> enable`.
local-next-hop enable	Enables the local next hop for this static route. The default form of this command is `default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> local-next-hop enable`. The no form of this command is `no ip route 0.0.0.0 0.0.0.0 <A.B.C.D> local-next-hop enable`.

Variable	Value
next-hop-vrf *WORD<1-16>*	Specifies the next-hop VRF instance by name. The default form of this command is `default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> next-hop-vrf WORD<1-16>`. The no form of this command is `no ip route 0.0.0.0 0.0.0.0 <A.B.C.D> next-hop-vrf WORD<1-16>`.
weight *<1-65535>*	Specifies the static route cost. The default form of this command is `default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> weight`.
preference *<1-255>*	Specifies the route preference. The default form of this command is `default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> preference`.

Enabling ICMP Router Discovery globally

Before you begin

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

About this task

Enable Router Discovery globally so that the device supports Router Discovery. Use ICMP Router Discovery to enable hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers.

If you enable ICMP Router Discovery globally, you automatically enable it for all VLANs. If you do not require ICMP Router Discovery on a specific VLAN, you must manually disable the feature.

Procedure

1. Enable ICMP Router Discovery on the device:

   ```
   ip irdp
   ```

2. Confirm that Router Discovery is enabled:

   ```
   show ip irdp [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Enable ICMP Router Discovery of the device:

```
VSP-9012:1(config)#ip irdp
```

confirm that Router Discovery is enabled:

```
VSP-9012:1(config)#show ip irdp
```

Variable definitions

Use the data in the following table to **show ip irdp** command.

Table 50: Variable definitions

Variable	Value
interface	Displays route discovery interface information.
vrf *WORD<1–16>*	Displays route discovery for particular VRF.
vrfids *WORD<0–512>*	Displays route discovery for particular VRF IDs.

Configuring Router Discovery on a port or VLAN

Before you begin

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

About this task

Enable Router Discovery so that the device forwards Router Discovery Advertisement packets to the VLAN or port.

Procedure

1. Specify the address placed in advertisement packets:

   ```
   ip irdp address <A.B.C.D>
   ```

2. Enable the interface to send the advertisement packets:

   ```
   ip irdp multicast
   ```

3. Configure other Router Discovery parameters for the interface as required.

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Log on to the GigabitEthernet Interface mode:

```
VSP-9012:1(config)#interface gigabitethernet 4/16
```

Specify the address placed in advertisement packets to the all-systems multicast address:

```
VSP-9012:1(config-if)#ip irdp address 244.0.0.1
```

Enable the interface to send the advertisement packets:

```
VSP-9012:1(config-if)#ip irdp multicast
```

Configure the lifetime for advertisements:

```
VSP-9012:1(config-if)#ip irdp holdtime 180
```

Variable definitions

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

Table 51: Variable definitions

Variable	Value
address <A.B.C.D>	Specifies the IP destination address use for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default address is 255.255.255.255. The default form of this command is `default ip irdp address`.
holdtime <4-9000>	Configures the lifetime for advertisements. The default form of this command is `default ip irdp holdtime`.
maxadvertinterval <4-1800>	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the router interface. The default is 600 seconds. The default form of this command is `default ip irdp maxadvertinterval`.
minadvertinterval <3-1800>	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to maxadvertinterval. The default is 450 seconds. The default form of this command is `default ip irdp minadvertinterval`.
multicast	Specifies if multicast advertisements are sent. The no form of this command is `no ip irdp multicast`.
preference <-2147483648-2147483647>	Specifies the preference (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet The default is 0. The default form of this command is `default ip irdp preference`.

Configuring a CLIP interface

Before you begin

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

About this task

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 your device. You can configure a maximum of 256 CLIP interfaces on each device.

Procedure

1. Create or access a CLIP interface:

   ```
   interface loopback <1-256>
   ```

   <1-256> indicates the identification number for the CLIP.

   The command prompt changes to indicate you now access the Loopback Interface Configuration mode.

2. Configure an IP address for the interface:

   ```
   ip address [<1-256>] <A.B.C.D/X> [vrf WORD<1-16>]
   ```

3. Enable OSPF on the CLIP interface:

   ```
   ip ospf [<1-256>] [vrf WORD<1-16>]
   ```

   You can configure other protocols on the CLIP interface; OSPF is the most common. See the following variable definitions table for other options.

4. View the IP address on the CLIP interface:

   ```
   show ip interface
   ```

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Create or access a CLIP interface:

```
VSP-9012:1(config)#interface loopback 200
```

Configure an IP address for the interface:

```
VSP-9012:1(config-if)#ip address 200 45.17.159.120/255.255.0.0
```

Enable OSPF on the CLIP interface:

```
VSP-9012:1(config-if)#ip ospf 200
```

View the IP address on the CLIP interface:

```
VSP-9012:1(config-if)#show ip interface
```

Variable definitions

Use the data in the following table to use the **ip** commands.

Table 52: Variable definitions

Variable	Value
address [<*1-256*>] <*A.B.C.D/X*> [vrf *WORD*<*1-16*>]	Specifies the IP address for the CLIP interface. *<1-256>* specifies the interface. <A.B.C.D/X> specifies the IP address and mask (0–32). vrf WORD<1-16> specifies an associated VRF by name. The no form of this command is `no ip address [<1-32>] <A.B.C.D> [vrf WORD<1-16>]`.
area <*1-256*> <*A.B.C.D*>[vrf *WORD*<*1-16*>]	Designates an area for the CLIP interface. <A.B.C.D> is the IP address of the OSPF area that is associated with the CLIP. vrf WORD<1-16> specifies an associated VRF by name The default form of this command is `default ip area <1-256> <A.B.C.D> [vrf WORD<1-16>]`. The no form of this command is `no ip area <1-256> vrf WORD<1-16>]`.
ospf [<*1-256*>] [vrf <*WORD 1-16*>]	Enables OSPF for the CLIP interface. *<1-256>* specifies the interface. vrf <WORD 1-16> specifies an associated VRF by name. The default form of this command is `default ip ospf <1-256> [vrf <WORD 1-16>]`. The no form of this command is `no ip ospf <1-256> [vrf <WORD 1-16>]`.
pim [<*1-256*>] [bsr-candidate preference <*0–255*>	Enables PIM for the CLIP interface. You can also enable the CLIP interface as a candidate bootstrap router and configure a preference value. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. The default form of this command is `default ip pim <1-256> [bsr-candidate]`. The no form of this command is `no ip pim <1-256> [bsr-candidate]`.

Chapter 11: IP routing configuration using Enterprise Device Manager

Configure the IP router interface so that you can use routing protocols and features on the interface. This section contains instructions for both the Global Router and Virtual Router Forwarding (VRF) instances.

Enabling routing for a router or a VRF instance

About this task

Enable IP forwarding (routing) on a router or a Virtual Router Forwarding (VRF) instance so that they support routing. You can use the IP address of any physical or virtual router interface for an IP-based network management.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Globals** tab.
4. To enable routing, select **Forwarding**.
5. Click **Apply**.

Enabling or disabling routing on a port

About this task

Enable or disable routing on a port to match your routing requirements. For example, you can disable routing on a particular port even if the port is part of a routed VLAN.

Procedure

1. From the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.

3. Click **General**.
4. Click the **Interface** tab.
5. In the Interface tab, select **enable** in the **AdminRouting** box to enable routing.

   OR

   In the Interface tab, select **disable** in the **AdminRouting** box to configure the port for bridging (and disable routing on this port).
6. Click **Apply**.

Deleting a dynamically-learned route

About this task

Use the Routes tab to view and manage the contents of the system routing table. You can also delete a dynamically learned route using this table. Exercise caution if you delete entries from the route table.

To delete a static route, use the **StaticRoute** tab.

To delete dynamic routes from the table for a VRF instance, first select the appropriate instance.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Routes** tab.
4. To delete a route, select the route and click **Delete**.

Routes field descriptions

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

Name	Description
Dest	Specifies the destination IP network of this route. An entry with a value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to multiple entries depends on the table access mechanisms defined by the network management protocol in use.
Mask	Indicates the network mask to logically add with the destination address before comparison to the destination IP network.
NextHop	Specifies the IP address of the next hop of this route.
NextHopId	Specifies the identifier of the next-hop, hostname or MAC address.

Name	Description
HopOrMetric	Specifies the primary routing metric for this route. The semantics of this metric are specific to various routing protocols.
Interface	Specifies the router interface for this route. • Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation. • Brouter interfaces are identified by the slot and port number of the brouter port.
Proto	Specifies the routing mechanism through which this route was learned: • other—none of the following • local—nonprotocol information, for example, manually configured entries • static • ICMP • EGP • GGP • Hello • RIP • IS-IS • ES-IS • Cisco IGRP • bbnSpfIgp • OSPF • BGP • Inter-VRF Redistributed Route
Age	Specifies the number of seconds since this route was last updated or otherwise determined correct.
PathType	Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths. • iA indicates Indirect Alternative route without an ECMP path • iAE indicates Indirect Alternative ECMP path • iB indicates Indirect Best route without ECMP path • iBE indicates Indirect Best ECMP path • dB indicates Direct Best route • iAN indicates Indirect Alternative route not in hardware • iAEN indicates Indirect Alternative ECMP route not in hardware

Name	Description
	• iBN indicates Indirect Best route not in hardware • iBEN indicates Indirect Best ECMP route not in hardware • dBN indicates Direct Best route not in hardware • iAU indicates Indirect Alternative Route Unresolved • iAEU indicates Indirect Alternative ECMP Unresolved • iBU indicates Indirect Best Route Unresolved • iBEU indicates Indirect Best ECMP Unresolved • dBU indicates Direct Best Route Unresolved • iBF indicates Indirect Best route replaced by FTN • iBEF indicates Indirect Best ECMP route replaced by FTN • iBV indicates Indirect best IPVPN route • iBEV indicates Indirect best ECMP IP VPN route • iBVN indicates Indirect best IP VPN route not in hardware • iBEVN indicates Indirect best ECMP IP VPN route not in hardware
Pref	Specifies the preference.
Layer3VirtualInterfaceType	Identifies the type for the value in the **Layer3VirtualInterface** field. The values include: • none – Specifies the type is not applicable for the route. • spb – Specifies that the routes are learned through ISIS and SPBM.
Layer3VirtualInterface	Specifies the layer 3 virtual interface. The values include: • 0 – Specifies the points for the SPB learned routes. • -1 – Specifies the route is not applicable.

Configuring IP route preferences

Before you begin

- Disable ECMP before you configure route preferences.

About this task

Change IP route preferences to force the routing protocols to prefer a route over another. Configure IP route preferences to override default route preferences and give preference to routes learned for a specific protocol.

Important:

Changing route preferences is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. Therefore, Avaya recommends that

if you want to change default preferences for routing protocols, do so before you enable the protocols.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **RoutePref** tab.
4. In the **Configured** column, change the preference for the given protocol.
5. Click **Apply**.

RoutePref field descriptions

Use the data in the following table to use the **RoutePref** tab.

Name	Description
Protocol	Specifies the protocol name.
Default	Specifies the default preference value for the specified protocol.
Configured	Configures the preference value for the specified protocol.

Flushing routing tables by VLAN

About this task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use Enterprise Device Manager (EDM) to flush the routing tables by VLAN or by port. Use this procedure to flush the IP routing table for a VLAN.

To flush routing tables by VLAN for a VRF instance, first select the appropriate instance.

Procedure

1. In the navigation tree, expand the following folders:**Configuration** > **VLAN**.
2. Click **VLANS**.
3. Click the **Advanced** tab.
4. In the **Vlan Operation Action** column, select a flush option.

   In a VLAN context, all entries associated with the VLAN are flushed. You can flush the ARP entries and IP routes for the VLAN.
5. Click **Apply**.

Flushing routing tables by port

About this task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use EDM to flush the routing tables by VLAN or flush them by port. Use this procedure to flush the IP routing table for a port.

To flush routing tables by port for a VRF instance, first select the appropriate instance.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
2. Click **General**.
3. Click the **Interface** tab.
4. In the **Action** section, select **flushAll**.

   In a port context, all entries associated with the port are flushed. You can flush the ARP entries and IP routes for a port.

   After you flush a routing table, it is not automatically repopulated. The repopulation time delay depends on the routing protocols in use.
5. Click **Apply**.

Assigning an IP address to a port

Before you begin

- Ensure routing (forwarding) is globally enabled.
- Ensure the VLAN is configured.
- If required, ensure the VRF instance exists.

About this task

Assign an IP address to a port so that it acts as a routable VLAN (a brouter port) and supports IP routing.

To configure a brouter port, assign an IP address to an IP policy-based single-port VLAN.

Important:

After you configure the IP address, you cannot edit the IP address, and you can assign only one IP address to any router interface (brouter or virtual).

You cannot assign an IP address to a brouter port that is a member of a routed VLAN. To assign an IP address to the brouter port, you must first remove the port from the routed VLAN.

If you want to assign a new IP address to a VLAN or brouter port that already has an IP address, first delete the existing IP address and then insert the new IP address.

Procedure

1. In Device Physical View, select the port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click **Insert**.
5. In the **Insert IP Address** dialog box, type the IP address, network mask, and VLAN ID.
6. Click **Insert**.

IP Address field descriptions

Use the data in the following table to help use the **IP Address** tab.

Name	Description
Interface	Specifies the router interface. • The name of the VLAN followed by the VLAN designation identifies virtual router interfaces. • The slot and port number of the brouter port identifies brouter interfaces.
Ip Address	Specifies the IP address of the brouter interface on this port. You can define only one IP address on a given port interface.
Net Mask	Specifies the subnet mask of the brouter interface on this port. The mask is an IP address with all the network bits set to 1 and all the hosts bits set to 0.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface.
ReasmMaxSize	Specifies the size of the largest IP packet which the interface can reassemble from fragmented incoming IP packets.
VlanId	Specifies the ID of the VLAN associated with the brouter port. This parameter is used to tag ports.
BrouterPort	Indicates whether this is a brouter port.
MacOffset	Translates the IP address into a MAC address. You can configure A MAC offset while you configure an IP address or the system can allot one within the allowed range.
VrfId	Specifies the associated VRF interface. The VrfId associates VLANs or brouter ports to a VRF after the creation of VLANs or brouter ports. VRF ID 0 is reserved for the Global Router.

Assigning an IP address to a VLAN

Before you begin

- Ensure routing (forwarding) is globally enabled.
- Ensure VLAN is configured.
- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

Specify an IP address for a VLAN so that the VLAN can perform IP routing.

Important:

You can assign only one IP address to any router interface (brouter or VLAN).

You cannot assign an IP address to a VLAN if a brouter port is a member of the VLAN. To assign an IP address to the VLAN, you must first remove the brouter port member.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP**.
5. Click **Insert**.
6. In the **Insert IP Address** dialog box, type the IP address and network mask.
7. Click **Insert**.

Viewing IP addresses for all router interfaces

About this task

Use the Addresses tab to view IP addresses (and their associated router interfaces) from one central location.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Addresses** tab.

Addresses field descriptions

Use the data in the following table to use the **Addresses** tab.

Name	Description
Interface	Specifies the router interface. • The name of the VLAN followed by the VLAN designation identifies virtual router interfaces. • The slot and port number of the brouter port identifies brouter interfaces.
Ip Address	Specifies the IP address of the router interface.
Net Mask	Specifies the subnet mask of the router interface.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface; that is, whether 0 (zero) or one is used for the broadcast address. Virtual Services Platform 9000 uses 1.
ReasmMaxSize	Specifies the size of the largest IP packet that this interface can reassemble from incoming fragmented IP packets.
VlanId	Identifies the VLAN associated with this entry. This value corresponds to the lower 12 bits in the IEEE 802.1Q VLAN tag.
BrouterPort	Indicates whether this is a brouter port (as opposed to a routable VLAN).
MacOffset	Specifies a number by which to offset the MAC address of the brouter port from the chassis MAC address. This ensures that each IP address has a different MAC address.

Configuring IP routing features globally

About this task

Configure the IP routing protocol stack to determine which routing features the Avaya Virtual Services Platform 9000 can use.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Globals** tab.
4. To globally enable routing, select **Forwarding**.
5. To globally configure the default TTL parameter type a value in the **DefaultTTL** field.

   This value is placed into routed packets that have no TTL specified.

6. To globally enable the Alternative Route feature, select **AlternativeEnable**.
7. To globally enable ICMP Router Discovery, select **RouteDiscoveryEnable**.
8. To globally enable ECMP, select **EcmpEnable**.
9. Configure the remaining parameters as required.
10. Click **Apply**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Forwarding	Configures the system for forwarding (routing) or nonforwarding. The default value is forwarding.
DefaultTTL	Configures the default time-to-live (TTL) value for a routed packet. TTL indicates the maximum number of seconds elapsed before a packet is discarded. Enter an integer from 1 to 255. The default value of 255 is used if a value is not supplied in the datagram header.
ReasmTimeout	Specifies the maximum number of seconds that received fragments are held while they wait for reassembly. The default value is 30 seconds.
ICMPUnreachableMsgEnable	Enables the generation of Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this system. These messages help determine if the system is reachable over the network. The default is disabled. **Important:** Avaya recommends that you only enable icmp-unreach-msg if it is absolutely required. If icmp-unreach-msg is enabled and a packet is received for which there is no route in the routing table, CPU utilization can dramatically increase.
ICMPRedirectMsgEnable	Enables or disables the system sending ICMP destination redirect messages.
AlternativeEnable	Globally enables or disables the Alternative Route feature.

Name	Description
	If the alternative-route parameter is disabled, all existing alternative routes are removed. After the parameter is enabled, all alternative routes are re-added. The default is enabled.
RouteDiscoveryEnable	Enables the ICMP Router Discovery feature. The default is disabled (not selected). Use ICMP Router Discovery to enable hosts attached to multicast or broadcast networks to discover the IP addresses of neighboring routers.
AllowMoreSpecificNonLocalRouteEnable	Enables or disables a more-specific nonlocal route. If enabled, the system can enter a more-specific nonlocal route into the routing table. The default is disabled.
SuperNetEnable	Enables or disables supernetting. If supernetting is globally enabled, the system can learn routes with a route mask less than 8 bits. Routes with a mask length less than 8 bits cannot have ECMP paths, even if you globally enable the ECMP feature. The default is disabled.
ARPLifeTime	Specifies the lifetime of an ARP entry within the system, global to Virtual Services Platform 9000. The default value is 360 minutes. The range for this value is 1 to 32767 minutes.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled. After ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.
EcmpMaxPath	Globally configures the maximum number of ECMP paths. • The interval is 1 to 8. • The default value is 1. You cannot configure this feature unless ECMP is enabled globally.
Ecmp1PathList	Selects a preconfigured ECMP path.
Ecmp2PathList	Selects a preconfigured ECMP path.
Ecmp3PathList	Selects a preconfigured ECMP path.
Ecmp4PathList	Selects a preconfigured ECMP path.
Ecmp5PathList	Selects a preconfigured ECMP path.
Ecmp6PathList	Selects a preconfigured ECMP path.

Name	Description
Ecmp7PathList	Selects a preconfigured ECMP path.
Ecmp8PathList	Selects a preconfigured ECMP path.
EcmpPathListApply	Applies changes in the ECMP path list configuration, or in the prefix lists configured as the path lists.

Configuring ECMP globally

About this task

Enable Equal Cost MultiPath (ECMP) to permit routers to determine up to eight equal-cost paths to the same destination prefix. You can use the multiple paths for load-sharing of traffic, which allows fast convergence to alternative paths. By maximizing load sharing among equal-cost paths, you can maximize the efficiency of your links between routers.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Select the **EcmpEnable** check box.
4. In the **EcmpMaxPath** box, enter the preferred number of equal-cost paths.
5. Click **Apply**.
6. Click **Close**.

Enabling alternative routes globally

Before you begin

- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

Globally enable alternative routes so that you can subsequently enable it on interfaces.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Select **AlternativeEnable**.

If the **AlternativeEnable** parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are re-added.

4. Click **Apply**.

Configuring static routes

About this task

Use static routes to force the router to make certain forwarding decisions. Create IP static routes to manually provide a path to destination IP address prefixes. The maximum number of static routes is 2000.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Static Routes** tab.
4. Click **Insert**.
5. If required, in the **OwnerVrfId** check box, select the appropriate VRF ID.
6. In the **Dest** field, type the IP address.
7. In the **Mask** field, type the subnet mask.
8. In the **NextHop** field, type the IP address of the router through which the specified route is accessible.
9. In the **NextHopVrfId** field, select the appropriate value.
10. To enable the static route, select the **Enable** check box.
11. In the **Metric** field, type the metric.
12. In the **Preference** field, type the route preference.
13. If required, select the **LocalNextHop** check box.

    Use this option to create Layer 3 static routes.
14. Click **Insert**.

    The new route appears in the **IP** dialog box, **Static Routes** tab.

Static Routes field descriptions

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

Name	Description
OwnerVrfId	Specifies the VRF ID for the static route.
Dest	Specifies the destination IP address of this route. A value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to such multiple entries is dependent on the table-access mechanisms defined by the network management protocol in use.
Mask	Indicates the mask that the system operates a logically AND function on, with the destination address, to compare the result to the Route Destination. For systems that do not support arbitrary subnet masks, an agent constructs the Route Mask by determining whether it belongs to a class A, B, or C network, and then uses one of: 255.0.0.0—Class A 255.255.0.0—Class B 255.255.255.0—Class C If the Route Destination is 0.0.0.0 (a default route) then the mask value is also 0.0.0.0.
NextHop	Specifies the IP address of the next hop of this route. In the case of a route bound to an interface which is realized through a broadcast media, the Next Hop is the IP address of the agent on that interface. When you create a black hole static route, configure this parameter to 255.255.255.255.
NextHopVrfId	Specifies the next-hop VRF ID in interVRF static route configurations. Identifies the VRF in which the ARP entry resides.
Enable	Determines whether the static route is available on the port. The default is enable. If a static route is disabled, it must be enabled before it can be added to the system routing table.
Status	Specifies the status of the route. The default is enabled.
Metric	Specifies the primary routing metric for this route. The semantics of this metric are determined by the routing protocol specified in the route RouteProto value. If this metric is not used, configure the value to 1. The default is 1.
IfIndex	Specifies the route index of the Next Hop. The interface index identifies the local interface through which the next hop of this route is reached.
Preference	Specifies the routing preference of the destination IP address. If more than one route can be used to forward IP traffic, the route that has the highest preference is used. The higher the number, the higher the preference.
LocalNextHop	Enables and disables LocalNextHop. If enabled, the static route becomes active only if the system has a local route to the network. If disabled, the static route becomes active if the system has a local route or a dynamic route.

Deleting a static route

About this task

Delete static routes that are no longer needed to prevent routing errors.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Static Routes** tab.
4. Select the static route you want to delete.
5. Click **Delete**.

Configuring a default static route

Before you begin

- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

The default route specifies a route to all networks for which there no explicit routes exist in the Forwarding Information Base or in the routing table. This route has a prefix length of zero (RFC 1812). You can configure Virtual Services Platform 9000 systems with a static default route, or they can learn it through a dynamic routing protocol.

To create a default static route, you configure the destination address and subnet mask to 0.0.0.0.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Static Routes** tab.
4. Click **Insert**.
5. In the **OwnerVrfId** check box, select the appropriate VRF ID.
6. In the **Dest** field, type 0.0.0.0.
7. In the **Mask** field, type 0.0.0.0.
8. In the **NextHop** field, type the IP address of the router through which the specified route is accessible.
9. In the **Metric** field, type the HopOrMetric value.
10. Click **Insert**.

Configuring a black hole static route

Before you begin

- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

Create a black hole static route to the destination that a router advertises to avoid routing loops when aggregating or injecting routes to other routers.

If an existing black hole route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Static Routes** tab.
4. Click **Insert**.
5. In the **OwnerVrfId** check box, select the appropriate VRF ID.
6. In the **Dest** field, enter the IP address.
7. In the **Mask** field, enter the network mask.
8. In the **NextHop** field, type 255.255.255.255.

   To create a black hole static route, you must configure the NextHop address to 255.255.255.255.
9. Select the **enable** option.
10. In the **Metric** box, type the HopOrMetric value.
11. In the **Preference** check box, select the route preference.

    When you specify a route preference, be sure to appropriately configure the preference so that when the black hole route is used, it is elected as the best route.
12. Click **Insert**.

Configuring ICMP Router Discovery globally

About this task

Enable ICMP Router Discovery so that it can operate on the system.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.

2. Click **IP**.
3. Click the **Globals** tab.
4. Select **RouteDiscoveryEnable**.
5. To select a preconfigured ECMP path, click the **EcmpPathList** ellipsis button.
6. Click **OK**.
7. Click **Apply**.
8. Click **Close**.

Configuring the ICMP Router Discovery table

Before you begin

- ICMP Router Discovery must be globally enabled.
- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

Configure the ICMP Router Discovery table to ensure correct ICMP operation for all interfaces that use Router Discovery.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Router Discovery** tab.
4. Configure the Router Discovery parameters to suit your network.
5. Click **Apply**.

Router Discovery field descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
Interface	Indicates the VLAN ID or the port.
AdvAddress	Specifies the IP destination address used for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address 224.0.0.1, or the limited-broadcast address 255.255.255.255. The default value is 255.255.255.255.

Name	Description
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface. The default value is true (advertise address).
AdvLifetime	Specifies the time to-live-value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 4 to 1800 seconds. The default value is 600 seconds.
MinAdvInterfal	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet. The range is –2147483648 to 2147483647. The default value is 0.

Configuring ICMP Router Discovery for a port

Before you begin

- You must globally enable ICMP Router Discovery.
- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

Use this procedure to configure Router Discovery on a port. When enabled, the port sends Router Discovery advertisement packets.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **Router Discovery** tab.
5. To enable Router Discovery, select **AdvFlag**.

6. Configure other parameters as required for proper operation.
7. Click **Apply**.

Router Discovery field descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
AdvAddress	Specifies the destination IP address used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address 224.0.0.1, or the limited-broadcast address 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface. The default value is True (advertise address).
AdvLifetime	Specifies the time to live value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 4 seconds to 1800 seconds. The default value is 600 seconds.
MinAdvInterval	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet. The accepted values are –2147483648 to 2147483647. The default value is 0.

Configuring ICMP Router Discovery on a VLAN

Before you begin

- You must globally enable ICMP Router Discovery.
- Change the VRF instance as required. For information about how to use EDM for a non0 VRF, see Selecting and launching a VRF context view on page 226.

About this task

Configure Router Discovery on a VLAN so that the ICMP Router Discovery feature can run over the VLAN. When enabled, the system sends Router Discovery advertisement packets to the VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select the VLAN ID that you want to configure to participate in Router Discovery.
4. Click **IP**.
5. Click the **Router Discovery** tab.
6. To enable Router Discovery for the VLAN, select **AdvFlag**.
7. Configure other parameters as required for proper operation.
8. Click **Apply**.

Router Discovery field descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
AdvAddress	Specifies the destination IP address used for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface. The default value is true (advertise address).
AdvLifetime	Specifies the time to-live-value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	Specifies the maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The range is 4 seconds to 1800 seconds. The default value is 600 seconds.
MinAdvInterval	The minimum time (in seconds) allowed between unsolicited broadcast or multicast router advertisements sent from the interface. The range is 3 seconds to MaxAdvInterval. The default value is 450 seconds.

Name	Description
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The range is –2147483648 to 2147483647. The default value is 0.

Configuring a CLIP interface

About this task

You can use a circuitless IP (CLIP) interface to provide uninterrupted connectivity to your system. You can configure a maximum of 256 CLIP interfaces on each device.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Circuitless IP** tab.
4. Click **Insert**.
5. In the **Interface** field, assign a CLIP interface number.
6. Enter the IP address.
7. Enter the network mask.
8. Click **Insert**.
9. To delete a CLIP interface, select the interface and click **Delete**.

Circuitless IP field descriptions

Use the data in the following table to use the **Circuitless IP** tab.

Name	Description
Interface	Specifies the number assigned to the interface, from 1 to 256.
Ip Address	Specifies the IP address of the CLIP.
Net Mask	Specifies the network mask.

Enabling OSPF on a CLIP interface

Before you begin

- You must globally enable OSPF.
- The OSPF area must already exist.

About this task

Enable Open Shortest Path First (OSPF) on a CLIP interface so that it can participate in OSPF routing.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Circuitless IP** tab.
4. Select the required CLIP interface.
5. Click **OSPF**.
6. Select the **Enable** check box.

   You must enable OSPF on the CLIP interface for CLIP to function.
7. In the current **AreaId** field, enter the IP address of the OSPF backbone area.
8. Click **Apply**.
9. Click **Close**.

Circuitless OSPF field descriptions

Use the data in the following table to use the **Circuitless OSPF** tab.

Name	Description
Enable	Enables OSPF on the CLIP interface.
AreaId	Specifies the OSPF area ID.

Enabling PIM on a CLIP interface

Before you begin

- You must globally enable PIM.

About this task

Enable Protocol Independent Multicasting (PIM) on a CLIP interface so that it can participate in PIM routing.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP**.
3. Click the **Circuitless IP** tab.
4. Select the required CLIP interface.
5. Click **PIM**.
6. Select the **Enable** check box.

   You must enable PIM on the CLIP interface for PIM to function. The mode is indicated on this tab.
7. Click **Apply**.
8. Click **Close**.

Circuitless PIM field descriptions

Use the descriptions in the following table to use the **Circuitless PIM** tab.

Name	Description
Enable	Enables PIM on the CLIP interface.
Mode	Specifies the PIM mode.

Viewing TCP global information

View TCP and UDP information to view the current configuration.

About this task

The fields on the TCP global tab provide information about the handshake (SYN) configuration and the maximum number of TCP connections you can create on your system.

When you initiate a TCP connection, both end points send handshake information to create the channel.

The retransmission algorithm and fields display the configured timeout value and minimum and maximum retransmission times that your system uses to terminate a connection attempt that falls outside your specified parameters.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP** or **Configuration** > **IPv6**.
2. Click **TCP/UDP**.

3. Click the **TCP Globals** tab.

TCP Global field descriptions

Use the data in the following table to use the **TCP Globals** tab.

Name	Description
RtoAlgorithm	Determines the timeout value used for retransmitting unacknowledged octets.
RtoMin	Displays the minimum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
RtoMax	Displays the maximum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
MaxConn	Displays the maximum connections for the device.

Viewing TCP connections information

View information about TCP connections.

About this task

Among other things, the fields on the TCP connections tab provide important information about the health of connections that traverse your switch.

In particular, the state column lets you know the state of each TCP connection. Of these, synSent, synReceived, and established indicate whether or not a channel is established and listen indicates when an end system is waiting for a returning handshake (SYN).

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP** or **Configuration** > **IPv6**.
2. Click **TCP/UDP**.
3. Click the **TCP Connections** tab.

TCP Connections field descriptions

Use the data in the following table to use the **TCP Connections** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
RemAddressType	Displays the type (IPv6 or IPv4) for the remote address of the TCP connection.
RemAddress	Displays the IPv6 address for the remote TCP connection.
RemPort	Displays the remote port number for the TCP connection.
State	Displays an integer that represents the state for the connection: • closed • listen • synSent • synReceived • established • finWait1 • finWait2 • closeWait • lastAck(9) • closing • timeWait • deleteTCB
Process	Displays the process ID for the system process associated with the TCP connection.

Viewing TCP listeners information

View TCP listener information.

About this task

The TCP listeners table provides a detailed list of systems that are in the listening state.

When a connection is in the listen state an end point system is waiting for a returning handshake (SYN).The normal listening state should be very transient, changing all of the time.

Two or more systems going to a common system in an extended listening state indicates the need for further investigation.

End systems in an extended listening state can indicate a broken TCP connection or a DOS attack on a resource.

This type of DOS attack, known as a SYN attack, results from the transmission of SYNs with no response to return replies.

While many systems can detect a SYN attack, the TCP listener statistics can provide additional forensic information.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP** or **Configuration** > **IPv6**.
2. Double-click **TCP/UDP**.
3. Click the **TCP Listeners** tab.

TCP Listeners field descriptions

Use the data in the following table to use the **TCP Listeners** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
Process	Displays the process ID for the system process associated with the TCP connection.

Chapter 12: RSMLT configuration using ACLI

Routed Split MultiLink Trunking (RSMLT) forwards packets in the event of core router failures, thus eliminating dropped packets during the routing protocol convergence.

Configuring RSMLT on a VLAN

Before you begin

- You must log on to the VLAN Interface Configuration mode in ACLI.
- You must enable the IP routing protocol on VLAN Layer 3 interfaces.
- VLANs with Layer 3 interfaces must also participate in Split MultiLink Trunking (SMLT).

About this task

Perform this procedure to configure RSMLT on each IP VLAN interface.

Use the no operator to disable RSMLT: `no ip rsmlt`

To configure this value to the default value, use the default operator with this command.

Procedure

Enable RSMLT on a VLAN:

```
ip rsmlt
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Log on to VLAN Interface Configuration mode:

```
VSP-9012:1(config)#interface VLAN 100
```

Enable RSMLT on a VLAN:

```
VSP-9012:1(config-if)#ip rsmlt
```

Variable definitions

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

Table 53: Variable definitions

Variable	Value
holddown-timer <0-3600>	Configures how long the RSMLT device does not participate in Layer 3 forwarding. *0-3600* is the timer value in seconds. To configure this value to the default value, use the default operator with this command. Avaya recommends that you configure this value to be longer than the anticipated routing protocol convergence.
holdup-timer <0-3600|9999>	Configures how long the RSMLT device maintains forwarding for its peer. *0-3600*|*9999* is the timer value in seconds. 9999 means infinity. To configure this value to the default value, use the default operator with this command.

Showing IP RSMLT information

Show IP RSMLT information to view data about all RSMLT interfaces.

Before you begin

Important:

If you use the `show ip rsmlt` command after you delete an RSMLT, the RSMLT still shows until you restart Avaya Virtual Services Platform 9000.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display RSMLT information about the interface:

   ```
   show ip rsmlt {edge-support] [<local|peer>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#show ip rsmlt local
```

```
================================================================================
                        Ip Rsmlt Local Info - GlobalRouter
================================================================================

VID   IP              MAC                ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
2     3.2.2.33        00:24:7f:9f:6a:04  Enable  Up    60     180
4     41.1.1.33       00:24:7f:9f:6a:08  Enable  Up    60     180

VID   SMLT ID
--------------------------------------------------------------------------------
2        1
4        1,  4

VID   IPv6            MAC                ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------

VID   SMLT ID
--------------------------------------------------------------------------------
```

Variable definitions

Use the information in the following command to use the **show ip rsmlt** command.

Table 54: Variable definitions

Variable	Value
edge-support	Displays the RSMLT edge-support and peer information
<local|peer>	Specifies values for the local or peer device.
vrf WORD<1-16>	Displays IP routing for a VRF.
vrfids WORD<0-512>	Displays IP routing for a range of VRFs.

Use the following table to use the **show ip rsmlt [<local|peer>]** command output.

Table 55: Variable definitions

Variable	Value
VID	Indicates the VLAN ID.
IP	Indicates the IP address of the router.
MAC	Indicates the MAC address assigned.
ADMIN	Indicates the administrative status of RSMLT on the router.
OPER	Indicates the operational status of RSMLT on the router.
HDTMR	Indicates the hold-down timer value in the range of 0 to 3600 seconds.
HUTMR	Indicates the hold-up timer value in the range of 0 to 3600 seconds or 9999. 9999 means infinity.
HDT REMAIN	Indicates the time remaining of the hold-down timer.
HUT REMAIN	Indicates the time remaining of the hold-up timer.

Variable	Value
SMLT ID	Indicates the Split MultiLink Trunk ID.

Configuring RSMLT edge support

Before you begin

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

Important:

If you use the `show ip rsmlt` command after you delete an RSMLT, the RSMLT still displays until you restart Avaya Virtual Services Platform 9000.

About this task

Configure RSMLT edge support to store the RSMLT peer MAC/IP address-pair in its local config file, and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT-peer systems. If enabled, all peer MAC/IP information for all RSMLT-enabled VLANs are saved during next the save config command.

RSMLT edge support is disabled by default.

Procedure

1. Enable RSMLT-edge:

   ```
   ip rsmlt edge-support
   ```

   Use the no operator to disable RSMLT-edge:`no ip rsmlt edge-support`

2. Clear RSMLT peer information, and then delete the RSMLT peer address:

   ```
   no ip rsmlt peer-address <1-4084>
   ```

3. Display RSMLT-edge status information:

   ```
   show ip rsmlt edge-support
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Enable RSMLT-edge:

```
VSP-9012:1(config)#ip rsmlt edge-support
```

Display RSMLT-edge status information:

```
VSP-9012:1(config)#show ip rsmlt edge-support
```

Variable definitions

Use the data in the following table to use the **`no ip rsmlt peer-address`** command.

Table 56: Variable definitions

Variable	Value
1–4084	Specifies the VLAN ID in the range of 1–4084.

Chapter 13: RSMLT configuration using Enterprise Device Manager

Routed Split MultiLink Trunking (RSMLT) forwards packets in the event of core router failures, thus eliminating dropped packets during the routing protocol convergence.

Configuring RSMLT on a VLAN

Configure RSMLT on a VLAN to exchange Layer 3 information between peer nodes in a switch cluster. This configuration applies to both IPv4 and IPv6.

Before you begin

- Enable an IP routing protocol on VLAN Layer 3 interfaces.
- Ensure VLANs with Layer 3 interfaces participate in Split MultiLink Trunking (SMLT).

About this task

The VLAN can be either IPv4 or IPv6, or both. RSMLT configuration on a VLAN simultaneously affects both IPv4 and IPv6.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **RSMLT** tab.
7. Select **Enable**.
8. In the **HoldDownTimer** field, type a hold-down timer value.
9. In the **HoldUpTimer** field, type a holdup timer value.
10. Click **Apply**.

RSMLT field descriptions

Use the data in the following table to use the **RSMLT** tab.

Name	Description
Enable	Enables RSMLT. The default is disabled.
HoldDownTimer	Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The range of this value is from 0 to 3600 seconds. The default is 60. If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 180. If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.

Viewing and editing RSMLT local information

About this task

Perform the following procedure to view and edit RSMLT local VLAN information.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **RSMLT**.
3. Click the **Local** tab.
4. Configure the parameters as required.
5. Click **Apply**.

Local field descriptions

Use the data in the following table to use the **Local** tab.

Name	Description
IfIndex	Specifies the IP route SMLT operation index.
VlanId	Specifies the VLAN ID of the chosen VLAN.

Name	Description
IpAddr	Specifies the IP address of the VLAN when RSMLT is enabled.
MacAddr	Specifies the MAC address of the selected VLAN.
Enable	Displays the RSMLT operating status as enabled or disabled.
OperStatus	Displays the RSMLT operating status as either up or down. The default is down.
HoldDownTimer	Defines how long the recovering/restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The range of this value is from 0 to 3600 seconds. The default is 0.
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 0.
SmltId	Specifies the ID range for the SMLT. A valid range is 1 to 512.
VrfId	Identifies the VRF.
VrfName	Indicates the VRF name.

Viewing RSMLT peer information

About this task

Perform this procedure to view and edit RSMLT peer information.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **RSMLT**.
3. Click the **Peer** tab.

Peer field descriptions

Use the following table to use the **Peer** tab.

Name	Description
IfIndex	Specifies the IP route SMLT operation index.
VlanId	Specifies the VLAN ID of the chosen VLAN.
IpAddr	Specifies the IP address of the VLAN when RSMLT is enabled.
MacAddr	Specifies the MAC address of the selected VLAN.
Enable	Displays the RSMLT operating status as enabled or disabled.
OperStatus	Displays the RSMLT operating status as either up or down. The default is down.

Name	Description
HoldDownTimer	Defines how long the recovering/restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The range of this value is from 0 to 3600 seconds. The default is 0.
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 0.
HoldDownTimeRemaining	Displays the time remaining of the HoldDownTimer. The default is 0.
HoldUpTimeRemaining	Displays the time remaining of the HoldUpTimer. The default is 0.
SmltId	Specifies the ID range for the Split MultiLink Trunk. A valid range is 1 to 32.
VrfId	Identifies the VRF.
VrfName	Indicates the VRF name.

Enabling RSMLT Edge support

Enable RSMLT Edge support to store the RSMLT peer MAC and IP address-pair in the local configuration file and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT peer systems.

The default is disabled.

About this task

RSMLT Edge support configuration applies to both IPv4 and IPv6.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **RSMLT**.
3. Click the **Globals** tab.
4. Select **EdgeSupportEnable**.
5. Click **Apply**.

Viewing RSMLT edge support information

About this task

View RSMLT edge support information to verify the RSMLT peer MAC/IP address-pair in its local configuration file and restore the configuration if the peer does not restore it after a simultaneous restart of both RSMLT-peer systems.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **RSMLT**.
3. Click the **Edge Peers** tab.

Edge Peers field descriptions

Use the data in the following table to use the **Edge Peers** tab fields.

Name	Description
VlanId	Specifies the VLAN ID of the chosen VLAN.
PeerIpAddress	Specifies the peer IP address.
PeerMacAddress	Specifies the peer MAC address.
PeerVrfId	Identifies the Peer VRF.
PeerVrfName	Specifies the Peer VRF name.

Chapter 14: VRRP configuration using ACLI

With the current implementation of virtual router redundancy protocol (VRRP), one active master switch exists for each IP subnet. All other VRRP interfaces in a network are in backup mode.

On Virtual Services Platform 9000, you cannot directly check or set the virtual IP address on the standby CPU module. To check or set the virtual IP address on the standby CPU, you must configure the virtual IP address on the master CPU, save it to the config.cfg file, and then copy that file to the standby CPU module.

If you have VRRP and IP routing protocols (for example, Open Shortest Path First [OSPF]) configured on the same IP physical interface, you cannot select the interface address as the VRRP virtual IP address (logical IP address). Use a separate dedicated IP address for VRRP.

To modify the behavior of the VRRP failover mechanism, use the hold-down timer to allow the router enough time to detect and update the OSPF or RIP routes. The timer delays the preemption of the master over the backup, when the master becomes available. The hold-down timer has a default value of 0 seconds. Avaya recommends that you configure all of your routers to the identical number of seconds for the hold-down timer. In addition, you can manually force the preemption of the master over the backup before the delay timer expires.

Virtual Router Redundancy Protocol (VRRP) 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 shared between two or more routers connecting the common subnet to the enterprise network.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

Important:

Virtual Services Platform 9000, when it acts as a VRRP master, does not reply to Simple Network Management Protocol (SNMP) Get requests to the VRRP virtual interface address. However, if Virtual Services Platform 9000, acts as a VRRP master, and receives SNMP Get requests to its physical IP address, then it does respond.

The Simple Network Management Protocol (SNMP) provides facilities to manage and monitor network resources. An SNMP manager and agent communicate through the SNMP protocol. The manager sends queries and the agent responds. An SNMP Get request is a message that requests the values of one or more objects.

When you use the fast advertisement interval option to configure a master and backup device, you must enable the fast advertisement interval option on both systems for VRRP to work correctly. If you configure one device with the regular advertisement interval, and the other device with the fast advertisement interval, it causes an unstable state and drops advertisements.

Important:

Ensure that Routed Split MultiLink Trunking (RSMLT) is not configured on the VLAN.

Configuring VRRP on a port or a VLAN

Configure VRRP on a port or a VLAN to forward packets to the virtual IP addresses associated with the virtual router and customize the VRRP configuration.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable

   configure terminal

   interface GigabitEthernet {slot/port[-slot/port][,...]}
   ```

2. Configure VRRP on a port:

   ```
   ip vrrp <1-255> enable
   ```

3. Configure a backup VRRP address:

   ```
   ip vrrp address <1-255> <A.B.C.D>
   ```

4. Show the global VRRP configuration:

   ```
   show ip vrrp
   ```

5. Display the backup VRRP address:

   ```
   show ip vrrp address
   ```

Example

Configure VRRP on a port. Configure a backup VRRP address. Enable the VRRP backup master, and show the global VRRP configuration.

```
VSP-switch:1>enable
VSP-switch:1#configure terminal
VSP-switch:1(config)#interface gigabitethernet 4/27
VSP-switch:1(config-if)#ip vrrp 27 enable
VSP-switch:1(config-if)#ip vrrp address 1 10.0.128.10
VSP-switch:1(config-if)#ip vrrp 1 backup-master enable
```

```
VSP-switch:1(config-if)#show ip vrrp
SP65:1(config)#show ip vrrp

================================================================================
                      VRRP Global Settings - GlobalRouter
================================================================================
ping-virtual-address : enabled
send-trap            : enabled

VSP-9012:1#show ip vrrp address

================================================================================
                    VRRP Info - GlobalRouter
================================================================================

VRRP ID  P/V  IP            MAC               STATE  CONTROL  PRIO  ADV
--------------------------------------------------------------------------------
1        2    10.0.128.10   00:00:5e:00:01:01 Init   Disabled 100   1

1 out of 1 Total Num of VRRP Address Entries displayed.


VRRP ID  P/V  MASTER   UP TIME            HLD DWN  CRITICAL IP(ENABLED)
--------------------------------------------------------------------------------
-------
1        2    0.0.0.0  0 day(s), 00:00:00 0        0.0.0.0   (No)

1 out of 1 Total Num of VRRP Address Entries displayed.


VRRP ID  P/V   BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)
--------------------------------------------------------------------------------
1        2      enable         up                  200       (NO)
1 out of 1 Total Num of VRRP Address Entries displayed.
```

Variable definitions

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

Table 57: Variable definitions

Variable	Value
1-255	Specifies the number of the VRRP to create or modify.
action {none|preempt}	Enables the choice option to manually override the hold-down timer and force preemption. You can configurenone|preempt to preempt the timer or configure it as none to allow the timer to keep working. To configure this option to the default value, use the default operator with this command.
address <1-255> <A.B.C.D>	Configures the VRRP virtual IP address. A.B.C.D is the IP address of the master VRRP.

Variable	Value
	Note: Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device. Use the no operator to remove the IP address of the VRRP physical interface:`no ip vrrp address <1-255> <A.B.C.D>`. To configure this option to the default value, use the default operator with this command.
adver-int <1-255>	Configures the the time interval between sending VRRP advertisement messages. The range is between 1 and 255 seconds. This value must be the same on all participating routers. The default is 1. To configure this option to the default value, use the default operator with this command.
backup-master enable	Enables the VRRP backup master. This option is supported only on Split MultiLink Trunking (SMLT) ports. Use the no operator to disable the VRRP backup master: `no ip vrrp <1-255> backup-master enable`. To configure this option to the default value, use the default operator with this command. **Important:** Do not enable backup master if you enable critical IP.
critical-ip-addr <A.B.C.D>	Configures the critical IP address for VRRP. A.B.C.D is the IP address on the local router, which is configured so that a change in its state causes a role device in the virtual router (for example, from master to backup in case the interface goes down).
critical-ip enable	Enables the critical IP address option. Use the no operator to disable the critical IP address option: `no ip vrrp <1-255> critical-ip enable`. To configure this option to the default value, use the default operator with this command. **Important:** Do not enable Critical IP if backup master is enabled.
enable	Enables VRRP on the port. Use the no operator to disable VRRP on the port: `no ip vrrp <1-255> enable`. To configure this option to the default value, use the default operator with this command.
fast-adv enable	Enables the Fast Advertisement Interval. The default is disabled.

Variable	Value
	Use the no operator to disable VRRP on the port: `no ip vrrp <1-255> fast-adv enable`. To configure this option to the default value, use the default operator with this command.
fast-adv-int <200-1000>	Configures the Fast Advertisement Interval, the time interval between sending VRRP advertisement messages. *200-1000* is the range in milliseconds, and must be the same on all participating routers. The default is 200. You must enter values in multiples of 200 milliseconds. To configure this option to the default value, use the default operator with this command.
holddown-timer <0-21600>	Configures the behavior of the VRRP failover mechanism by allowing the router enough time to detect the OSPF or RIP routes. *0-21600* is the time interval (in seconds) a router is delayed when changing to master state. To configure this option to the default value, use the default operator with this command.
priority <1-255>	Configures the port VRRP priority. *1-255* is the value used by the VRRP router. The default is 100. Assign the value 255 to the router that owns the IP address associated with the virtual router. To configure this option to the default value, use the default operator with this command.

Showing VRRP port or VLAN information

Before you begin

- You must log on to the Privileged EXEC mode.

About this task

Show VRRP port or VLAN information to view configuration details and operational status.

Procedure

Display basic VRRP configuration information about the specified port, all ports, or the VLAN:

```
show ip vrrp address [vrid <1-255>] [addr <A.B.C.D>] [vrf WORD<1-16>]
[vrfids WORD<0-512>]
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#show ip vrrp address

================================================================================
```

```
                              VRRP Info - GlobalRouter
================================================================================

VRRP ID  P/V   IP              MAC                STATE    CONTROL  PRIO  ADV
--------------------------------------------------------------------------------
74       74    74.74.74.74     00:00:5e:00:01:4a  Init     Disabled 100   1
75       75    75.75.75.75     00:00:5e:00:01:4b  Init     Disabled 100   1
76       76    76.76.76.76     00:00:5e:00:01:4c  Init     Disabled 100   1
77       77    77.77.77.77     00:00:5e:00:01:4d  Init     Disabled 100   1
78       78    78.78.78.78     00:00:5e:00:01:4e  Init     Disabled 100   1
79       79    79.79.79.79     00:00:5e:00:01:4f  Init     Disabled 100   1
80       80    80.80.80.80     00:00:5e:00:01:50  Init     Disabled 100   1
81       81    81.81.81.81     00:00:5e:00:01:51  Init     Disabled 100   1
82       82    82.82.82.82     00:00:5e:00:01:52  Init     Disabled 100   1
83       83    83.83.83.83     00:00:5e:00:01:53  Init     Disabled 100   1
84       84    84.84.84.84     00:00:5e:00:01:54  Init     Disabled 100   1
85       85    85.85.85.85     00:00:5e:00:01:55  Init     Disabled 100   1
86       86    86.86.86.86     00:00:5e:00:01:56  Init     Disabled 100   1
87       87    87.87.87.87     00:00:5e:00:01:57  Init     Disabled 100   1
88       88    88.88.88.88     00:00:5e:00:01:58  Init     Disabled 100   1
89       89    89.89.89.89     00:00:5e:00:01:59  Init     Disabled 100   1
90       90    90.90.90.90     00:00:5e:00:01:5a  Init     Disabled 100   1
91       91    91.91.91.91     00:00:5e:00:01:5b  Init     Disabled 100   1
92       92    92.92.92.92     00:00:5e:00:01:5c  Init     Disabled 100   1
93       93    93.93.93.93     00:00:5e:00:01:5d  Init     Disabled 100   1

20 out of 20 Total Num of VRRP Address Entries displayed


VRRP ID  P/V   MASTER          UP TIME            HLD DWN  CRITICAL IP(ENABLED)
--------------------------------------------------------------------------------
74       74    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
75       75    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
76       76    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
77       77    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
78       78    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
79       79    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
80       80    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
81       81    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
82       82    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
83       83    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
84       84    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
85       85    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
86       86    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
87       87    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
88       88    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
89       89    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
90       90    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
91       91    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
92       92    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)
93       93    0.0.0.0         0 day(s), 00:00:00    0      0.0.0.0   (No)

20 out of 20 Total Num of VRRP Address Entries displayed


VRRP ID  P/V   BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)
--------------------------------------------------------------------------------
74       74    disable         down                  200       (NO)
75       75    disable         down                  200       (NO)
76       76    disable         down                  200       (NO)
77       77    disable         down                  200       (NO)
78       78    disable         down                  200       (NO)
79       79    disable         down                  200       (NO)
80       80    disable         down                  200       (NO)
81       81    disable         down                  200       (NO)
```

```
82        82    disable         down                  200        (NO)
83        83    disable         down                  200        (NO)
84        84    disable         down                  200        (NO)
85        85    disable         down                  200        (NO)
86        86    disable         down                  200        (NO)
87        87    disable         down                  200        (NO)
88        88    disable         down                  200        (NO)
89        89    disable         down                  200        (NO)
90        90    disable         down                  200        (NO)
91        91    disable         down                  200        (NO)
92        92    disable         down                  200        (NO)
93        93    disable         down                  200        (NO)

20 out of 20 Total Num of VRRP Address Entries displayed
```

Variable definitions

Use the data in the following table to use the **show ip vrrp address** command.

Table 58: Variable definitions

Variable	Value
addr <A.B.C.D>	Specifies the physical local address of the master VRRP.
vrf WORD<1–16>	Specifies the name of the VRF.
vrid <1-255>	Specifies a unique integer value that represents the virtual router ID in the range 1–255. The virtual router acts as the default router for one or more assigned addresses.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.

Use the data in the following table to use the **show ip vrrp address** command output.

Table 59: Variable definitions

Variable	Value
ADV	Indicates the Advertisement Interval, in milliseconds, between sending advertisement messages.
BACKUP MASTER	Indicates the backup master IP address.
BACKUP MASTER STATE	Indicates the backup master state.
CONTROL	Indicates the virtual router function. Configure the value to enabled to transition the state of the router from initialize to backup. Configure the value to disabled to transition the router from master or backup to initialize.
CRITICAL IP	Indicates the IP address of the interface that causes a shutdown event.
CRITICAL IP (ENABLED)	Indicates if the critical IP address is enabled.

Variable	Value
FAST ADV	Indicates the Fast Advertisement Interval, in milliseconds, between sending advertisement messages. When the Fast Advertisement Interval is enabled, the Fast Advertisement Interval is used instead of the regular advertisement interval.
FAST ADV (ENABLED)	Indicates the state of fast advertisement.
HLD DWN	Indicates the amount of time (in seconds) remaining until the Hold Down timer expires. If the Hold Down timer is not active, the value displays as 0.
IP	Indicates the assigned IP addresses that a virtual router backs up.
MAC	Indicates the virtual MAC address of the virtual router in the format 00-00-5E-00-01-<vrrpid>, where the first three octets consist of the IANA OUI; the next two octets indicate the address block of the VRRP protocol; and the remaining octets consist of the vrrpid.
MASTER	Indicates the master router real (primary) IP address. The master route is the IP address listed as the source in the VRRP advertisement last received by this virtual router.
PRIO	Indicates the priority for the virtual router (for example, master election) with respect to other virtual routers that are backing up one or more associated IP addresses. Higher values indicate higher priority. A priority of 0, which you cannot configure, indicates that this router does not participate in VRRP and a backup virtual router transitions to become the new master. A priority of 255 is used for the router that owns the associated IP addresses.
P/V	Indicates whether this device responds to pings directed to a IP address.
STATE	Indicates the current state of the virtual router. initialize—waiting for a startup event backup—monitoring the state or availability of the master router master—forwarding IP addresses associated with this virtual router.
UP TIME	Indicates the time interval (in hundredths of a second) since this virtual router was initialized.
VRID	Indicates the virtual router ID on a VRRP router.

Showing extended VLAN VRRP

Perform this procedure to display the extended VRRP configuration for all VLANs or a specified VLAN on the device.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Show the extended VRRP configuration for all VLANs on the device or for the specified VLAN:

   ```
   show ip vrrp interface vlan [<1-4084>] [verbose] [vrf WORD<1-16>]
   [vrfids WORD<0-512>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#show ip vrrp interface vlan
================================================================================
                                  Vlan Vrrp
================================================================================
VLAN VRF              VRRP IP              VIRTUAL
ID   NAME             ID   ADDRESS         MAC ADDRESS
--------------------------------------------------------------------------------
3    GlobalRouter     30   30.30.30.123    00:00:5e:00:01:1e

All 1 out of 1 Total Num of Vlan Vrrp displayed
```

Variable definitions

Use the data in the following table to use the **show ip vrrp interface vlan** command.

Table 60: Variable definitions

Variable	Value
<1-40984>	Specifies the VLAN ID.
portList	Specifies the slot or port number of a range of ports.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.

Use the data in the following table to use the **show ip vrrp interface vlan [<1-4084>] [verbose] [vrf WORD<1-16>] [vrfids WORD<0-512>]** command output.

Table 61: Variable definitions

Variable	Value
VLAN ID	Indicates the VLAN ID.
STATE	Indicates the current state of the virtual router. • initialize—waiting for a startup event • backup—monitoring the state or availability of the master router

Variable	Value
	• master—forwarding IP addresses associated with this virtual router
CONTROL	Indicates the virtual router function. Configure the value to enabled to transition the state of the router from initialize to backup. Configure the value to disabled to transition the router from master or backup to initialize.
PRIORITY	Indicates the priority for the virtual router (for example, master election) with respect to other virtual routers that are backing up one or more associated IP addresses. Higher values indicate higher priority. A priority of 0, which you cannot configure, indicates that this router ceased to participate in VRRP and a backup virtual router transitions to become a new master. Use a priority of 255 for the router that owns the associated IP addresses.
MASTER IPDDR	Indicates the master router real (primary) IP address. The master IP address is listed as the source in the VRRP advertisement last received by this virtual router.
ADVERTISE INTERVAL	Indicates the time interval, in seconds, between sending advertisement messages. Only the master router sends VRRP advertisements.
CRITICAL IPADDR	Indicates the IP address of the interface that causes a shutdown event.
HOLDDOWN_TIME	Indicates the configured time (in seconds) that the system waits before it preempts the current VRRP master.
ACTION	Indicates the trigger for an action on this VRRP interface. Options include none and preemptHoldDownTimer.
CRITICAL IP ENABLE	Indicates that a user-defined critical IP address is enabled. No indicates the use of the default IP address (0.0.0.0).
BACKUP MASTER	Indicates the state of designating a backup master router.
BACKUP MASTER STATE	Indicates the state of the backup master router.
FAST ADV INTERVAL	Indicates the time interval, in milliseconds, between sending Fast Advertisement messages. When the Fast Advertisement Interval is enabled, the Fast Advertisement Interval is used instead of the regular advertisement interval.
FAST ADV ENABLE	Indicates the Fast Advertisement Interval status.

Showing VRRP interface information

About this task

If you enter a virtual router ID or an IP address when showing VRRP interface information, the information appears only for that virtual router ID or for that interface.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display VRRP information about the interface:

   ```
   show ip vrrp interface [gigabitethernet {slot/port[-slot/port][,...]}][statistics][verbose][vlan <1-4084>] [vrfWORD<1-16>][vrfids WORD<0-512>][vrid<1-255>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1(config)#show ip vrrp interface
================================================================================
                                  Vlan Vrrp
================================================================================
VLAN VRF              VRRP IP              VIRTUAL
ID   NAME             ID   ADDRESS         MAC ADDRESS
--------------------------------------------------------------------------------
3    GlobalRouter     30   32.30.30.122    00:00:5e:00:01:1f

All 1 out of 1 Total Num of Vlan Vrrp displayed

================================================================================
                                  Port Vrrp
================================================================================
PORT  VRF              VRRP IP              VIRTUAL
NUM   NAME             ID   ADDRESS         MAC ADDRESS
--------------------------------------------------------------------------------
3/5    Router25               32.12.44.10               00:00:5e:00:01:1a
```

Variable definitions

Use the data in the following table to use the **show ip vrrp interface** command.

Table 62: Variable definitions

Variable	Value
gigabitethernet {slot/port[-slot/port][,...]}	Specifies to show the VRRP information of which interface.
statistics	Specifies VRRP statistics.
verbose	Specifies to show all available information about the VRRP interfaces.
vlan	Specifies the VLAN that contains the VRRP.
vrf WORD<1-16>	Specifies the name of the VRF.
vrid <1-255>	Specifies a unique integer value that represents the virtual router ID in the range 1–255. The virtual router acts as the default router for one or more assigned addresses.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.

Configuring VRRP notification control

Use the following procedure to enable VRRP notification control. The generation of SNMP traps for VRRP events is enabled, by default.

About this task

You can configure traps by creating SNMPv3 trap notifications, creating a target address to send the notifications, and specify target parameters. For more information about how to configure trap notifications, see

Troubleshooting Avaya Virtual Services Platform 9000, NN46250-700.

Procedure

1. Enter VRRP Router Configuration mode:

   ```
   enable
   configure terminal
   router vrrp
   ```

2. Enable a trap for VRRP events:

   ```
   send-trap enable [vrf WORD<0-16>]
   ```

3. Disable a trap for VRRP events:

   ```
   no send-trap enable [vrf WORD<0-16>]
   ```

4. Configure a trap for VRRP events to the default:

   ```
   default send-trap enable [vrf WORD<0-16>]
   ```

5. Display the configuration:

   ```
   show ip vrrp [vrf WORD<0-16>]
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrrp
Switch:1(config-vrrp)#send-trap enable vrf mgmtrouter
Switch:1(config)#show ip vrrp vrf mgmtrouter

====================================
VRRP Global Settings - VRF mgmtrouter
====================================
ping-virtual-address : enabled
send-trap : enabled
```

Variable definitions

Use the data in the following table to use the `send-trap` and `show ip vrrp` commands.

Variable	Value
enable	Enables generation of SNMP traps.
vrf *WORD<0–16>*	Configures the send-trap for a particular VRF.

Enabling ping to a virtual IP address

Use the following procedure to enable ping to a virtual IP address. The default is enabled.

Procedure

1. Enter VRRP Router Configuration mode:

   ```
   enable
   configure terminal
   router vrrp
   ```

2. Enable ping to a virtual IP address:

   ```
   ping-virtual—address enable [vrf WORD<0–16>]
   default ping-virtual—address enable [vrf WORD<0–16>]
   ```

3. Disable ping to a virtual IP address:

   ```
   no ping-virtual—address enable [vrf WORD<0–16>]
   ```

4. Display the configuration:

   ```
   show ip vrrp [vrf WORD<0–16>]
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrrp
Switch:1(config-vrrp)#ping-virtual-address enable vrf mgmtrouter
Switch:1(config)#show ip vrrp vrf mgmtrouter

====================================
VRRP Global Settings - VRF mgmtrouter
====================================
ping-virtual-address : enabled
send-trap : enabled
```

Variable definitions

Use the data in the following table to use the `ping-virtual—address enable` and `show ip vrrp` commands.

Variable	Value
enable	Enables ping to a virtual IP address.
vrf *WORD<0–16>*	Specifies the VRF.

Chapter 15: VRRP configuration using EDM

With the current implementation of Virtual Router Redundancy Protocol (VRRP), one active master switch exists for each IP subnet. All other VRRP interfaces in a network are in backup mode.

On the Avaya Virtual Services Platform 9000, you cannot directly check or set the virtual IP address on the standby CPU module. To check or set the virtual IP address on the standby CPU, you must configure the virtual IP address on the master CPU, save it to the config.cfg file, and then copy that file to the standby CPU module.

If you have VRRP and IP routing protocols (for example, Open Shortest Path First [OSPF]) configured on the same IP physical interface, you cannot select the interface address as the VRRP virtual IP address (logical IP address). Use a separate dedicated IP address for VRRP.

To modify the behavior of the VRRP failover mechanism, use the hold-down timer to allow the router enough time to detect and update the OSPF or RIP routes. The timer delays the preemption of the master over the backup, when the master becomes available. The hold-down timer has a default value of 0 seconds. Avaya recommends that you configure all of your routers to the identical number of seconds for the hold-down timer. In addition, you can manually force the preemption of the master over the backup before the delay timer expires.

Virtual Router Redundancy Protocol (VRRP) 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 shared between two or more routers connecting the common subnet to the enterprise network.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

Important:

The Avaya Virtual Services Platform 9000, when it acts as a VRRP master, does not reply to Simple Network Management Protocol (SNMP) Get requests to the VRRP virtual interface address. However, if Avaya Virtual Services Platform 9000, acts as a VRRP master, and receives SNMP Get requests to its physical IP address, then it does respond.

The Simple Network Management Protocol (SNMP) provides facilities to manage and monitor network resources. An SNMP manager and agent communicate through the SNMP protocol. The manager sends queries and the agent responds. An SNMP Get request is a message that requests the values of one or more objects.

When you use the fast advertisement interval option to configure a master and backup device, you must enable the fast advertisement interval option on both systems for VRRP to work correctly. If you configure one device with the regular advertisement interval, and the other device with the fast advertisement interval, it causes an unstable state and drops advertisements.

Important:

Ensure that Routed Split MultiLink Trunking (RSMLT) is not configured on the VLAN.

Before you begin

- Assign an IP address to the interface.
- Enable VRRP globally.
- Ensure RSMLT is not configured on the VLAN.

Enabling VRRP global variables

About this task

Enable VRRP global variables to enable the VRRP function.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRRP**.
3. Click the **Globals** tab.
4. Configure the required features.
5. Click **Apply**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
NotificationCntl	Indicates whether the VRRP-enabled router generates SNMP traps for events. • enabled—SNMP traps are generated • disabled—no SNMP traps are sent The default is enabled.
PingVirtualAddrEnable	Configures whether this device responds to pings directed to a virtual router IP address. The default is enabled.

Configuring VRRP for the interface

About this task

You can manage and configure VRRP parameters for the routing interface.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRRP**.
3. Click the **Interface** tab.
4. Double-click the **HoldDownTimer** field, and enter the number of seconds for the timer.

   The **HoldDownState** field displays active when the hold-down timer is counting down and preemption occurs. The field displays dormant when preemption is not pending. When the hold-down timer is active, the **HoldDownTimeRemaining** field displays the seconds remaining before preemption.
5. In the **Action** check box, select an option.
6. Click **Apply**.

Interface field descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Specifies the index value that uniquely identifies the interface to which this entry is applicable.
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
IpAddr	Specifies the assigned IP addresses that a virtual router is responsible for backing up.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Specifies the state of the virtual router interface: • Initialize—waiting for a startup event • Backup—monitoring availability and state of the master router • Master—functioning as the forwarding router for the virtual router IP addresses.

Name	Description
Control	Specifies whether VRRP is enabled or disabled for the port (or VLAN). The default is enabled.
Priority	Specifies the priority value used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvertisementInterval	Specifies the time interval (in seconds) between sending advertisement messages. The range is 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements. The default is 1.
MasterIpAddr	Specifies the IP address of the physical interface of the master virtual router that forwards packets sent to the virtual IP addresses associated with the virtual router.
VirtualRouterUpTime	Specifies the time interval (in hundredths of a second) since the virtual router was initialized.
Action	Lists options to override the delay timer manually and force preemption: • **none** does not override the timer • **preemptHoldDownTimer** preempts the timer
HoldDownTimer	Configures the amount of time (in seconds) to wait before preempting the current VRRP master.
HoldDownState	Indicates the hold-down state of this VRRP interface. If the hold-down timer is operational, this variable is set to active; otherwise, this variable is set to dormant.
HoldDownTimeRemaining	Indicates the amount of time (in seconds) left before the HoldDownTimer expires.
CriticalIpAddr	Specifies an IP interface on the local router configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.
CriticalIpAddrEnable	Configures the IP interface on the local router to enable or disable the backup. The default is disabled.
BackUpMaster	Enables the backup VRRP system traffic forwarding. This reduces the traffic on the IST link. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP system traffic forwarding is enabled or disabled. The default is disabled.
FasterAdvInterval	Configures the Fast Advertisement Interval between sending VRRP advertisement messages. The interval is between 200 and 1000 milliseconds, and you must enter the same value on all participating routers. The default is 200. You must enter the values in multiples of 200 milliseconds.
FasterAdvIntervalEnable	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disable.

Configuring VRRP on a port or a VRF instance

About this task

You can configure VRRP on a port, a brouter port (or a VLAN), or a VRF instance only if the port or brouter port (or VLAN) is assigned an IP address.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **VRRP** tab.
5. Click **Insert**.
6. In the **VrId** field, enter a virtual router ID.
7. In the **IpAddr** field, type an IP address.
8. In the **Control** section, select **enabled**.
9. In the **AdvertisementInterval** field, enter an advertisement interval.
10. If required, select an **Action**.
11. Specify the number of seconds for the HoldDown timer.
12. Enter a critical IP address.
13. Select **CriticalIpAddrEnable**.

    Important:

    Do not enable Critical IP if Backup Master is enabled.
14. In **BackUpMaster**, select **enabled**.
15. Click **Insert**.
16. Click **Close**.

VRRP field descriptions

Use the data in the following table to use the **VRRP** tab.

Name	Description
IfIndex	Specifies the port interface index.

Name	Description
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
IpAddr	Specifies the IP address of the virtual router interface.
VirtualMacAddr	Indicates the virtual MAC address of the virtual router using the following format: 00-00-5E-00-01-<virtual router ID>.
State	Shows the current state of the virtual router.
Control	Displays whether VRRP is enabled or disabled for the port or VLAN.
Priority	Specifies a priority value used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
FasterAdvIntervalEnable	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disabled.
FasterAdvInterval	Configures the Fast Advertisement Interval between sending VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. You must enter values in multiples of 200 milliseconds.
AdvertisementInterval	Specifies the time interval (in seconds) between sending advertisement messages. The range is 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements.
MasterIpAddr	Indicates the primary IP address of the master router. This IP address is the source in VRRP advertisements last received by the virtual router.
VirtualRouterUpTime	Indicates the time interval, in hundredths of a second, since the virtual router initialized.
Action	Use the action list to manually override the delay timer and force preemption: • **none** does not override the timer • **preemptHoldDownTimer** preempts the timer
HoldDownTimer	Specifies the time interval (in seconds) a router is delayed for the following conditions: • The VRRP hold-down timer runs when the system transitions from initialization to backup to master. This occurs only on a system startup. • The VRRP hold-down timer does not run under the following condition: In a nonstartup condition, the backup system becomes master after the Master Downtime

Name	Description
	Interval (3 * hello interval), if the master virtual router goes down. • The VRRP hold-down timer also applies to the VRRP BackupMaster feature.
HoldDownState	Indicates the hold-down state of the VRRP interface. If the hold-down timer is operational, this value is active.
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the hold-down timer expires.
CriticalIpAddr	Indicates if a user-defined critical IP address must be enabled. There is no effect if a user-defined IP address does not exist. Use the default IP address (0.0.0.0).
CriticalIpAddrEnable	Configures the IP interface on the local router to enable or disable the backup.
BackUpMaster	Enables the VRRP backup master feature. This option is only supported on SMLT ports. The default is disabled.
BackUpMasterState	Indicates if the backup VRRP switch is enabled for traffic forwarding. The default is disabled.

Configuring VRRP on a VLAN (or brouter port) or a VRF instance

Before you begin

- You must first configure VRRP globally before you configure VRRP on a VLAN or a VRF instance.
- You must assign an IP address to the port or VLAN before you can configure VRRP on a VLAN or brouter port.

About this task

Perform this procedure to configure VRRP on a VLAN, brouter port, or a VRF instance.

Note:

Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. In the **Basic** tab, select a VLAN.
4. Click **IP**.

5. Click the **VRRP** tab.
6. Click **Insert**.
7. Configure the VRRP feature as required.
8. Click **Insert**.

VRRP field descriptions

Use the data in the following table to use the **VRRP** tab.

Name	Description
IfIndex	Specifies the index value that uniquely identifies the interface to which this entry is applicable.
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
IpAddr	Specifies the IP address of the virtual router interface.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Specifies he state of the virtual router interface: • initialize—waiting for a startup event • backup—monitoring availability and state of the master router • master—functioning as the forwarding router for the virtual router IP addresses.
Control	Displays whether VRRP is enabled or disabled for the port or VLAN.
Priority	Specifies a priority value used by this VRRP router. The range is from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
FasterAdvIntervalEnable	Enables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disabled.
FasterAdvInterval	Configures the Fast Advertisement Interval between sending VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.
AdvertisementInterval	Specifies the time interval (in seconds) between sending advertisement messages. The range is from 1

Name	Description
	to 255 seconds with a default of 1 second. Only the master router sends advertisements.
MasterIPAddr	Specifies the IP address of the master router.
VirtualRouterUpTime	Specifies the time interval (in hundredths of a second) since the virtual router was initialized.
Action	Use the action list to manually override the delay timer and force preemption: • preemptHoldDownTimer—preempt the timer • none—allow the timer to keep working
HoldDownTimer	Specifies the time interval (in seconds) a router is delayed for the following conditions: • The VRRP hold-down timer runs when the system transitions from initialization to backup to master. This occurs only on a system startup. • The VRRP hold-down timer does not run under the following condition: In a nonstartup condition, the backup system becomes master after the Master Downtime Interval (3 * hello interval), if the master virtual router goes down. • The VRRP hold-down timer also applies to the VRRP BackupMaster feature.
HoldDownState	Configures the hold down state. The Status is active when the hold-down timer is counting down and preemption occurs; the status displays dormant when preemption is not pending.
HoldDownTimeRemaining	Specifies the seconds remaining before preemption.
CriticalIpAddr	Indicates if a user-defined critical IP address must be enabled. There is no effect if a user-defined IP address does not exist.
CriticalIpAddrEnable	Configures the IP interface on the local router to enable or disable the backup.
BackUpMaster	Enables the VRRP backup master feature. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP system traffic forwarding is enabled or disabled. The default is disabled.

Configuring Fast Advertisement Interval on a port or a VRF instance

About this task

Configure the Fast Advertisement Interval to send VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **VRRP** tab.
5. Click **Insert**.
6. In the**Insert VRRP** dialog box, enable **FasterAdvIntervalEnable**.
7. In the **FasterAdvInterval** field, enter a value. You must set this value using multiples of 200 milliseconds.
8. Click **Insert**. The new entry appears in the **VRRP** tab of the **Port** dialog box.

Configuring Fast Advertisement Interval on a VLAN or a VRF instance

About this task

Configure the Fast Advertisement Interval to send VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Select a VLAN.
4. Click **IP**.
5. Click the **VRRP** tab.
6. Click **Insert**.
7. In the IP, VLAN, Insert VRRP dialog box, click the **FasterAdvIntervalEnable** enable option.

8. In the **FasterAdvInterval**, box, enter a value. You must set the value using multiples of 200 milliseconds.
9. Click **Insert**. The new entry appears in the VRRP tab of the IP, VLAN dialog box.

Chapter 16: VRF Lite fundamentals

Use the concepts described in this section to understand and use the Virtual Routing and Forwarding (VRF) Lite feature. Use VRF Lite to provide secure customer data isolation.

Overview

Use VRF Lite to offer networking capabilities and traffic isolation to customers that operate over the same node (router). Each virtual router emulates the behavior of a dedicated hardware router; the network treats each virtual router as a separate physical router. In effect, you can perform the functions of many routers using a single platform that runs VRF Lite. The result is a substantial reduction in the cost associated with providing routing and traffic isolation for multiple clients.

The following figure shows one platform acting as multiple virtual routers, each serving a different customer network.

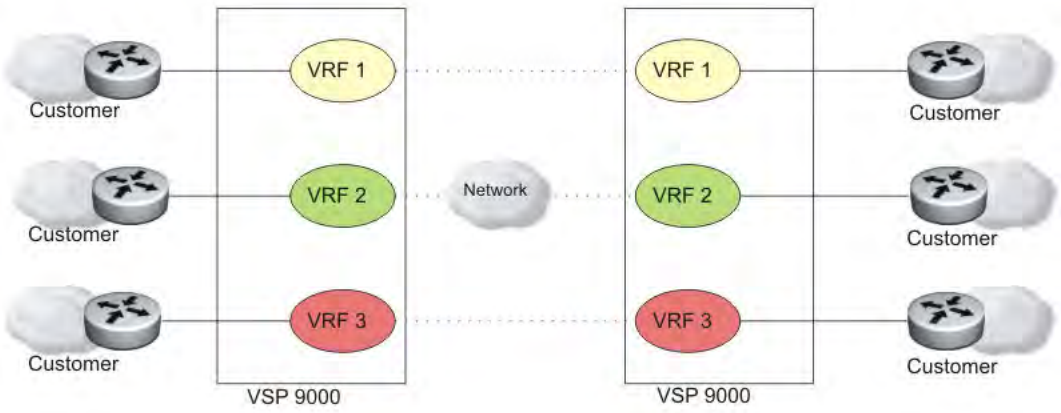


Figure 16: Multiple virtual routers in one system

One Avaya Virtual Services Platform 9000 can support many virtual routers. Each virtual router instance is called a VRF instance. A VRF represents a single instance of a virtual router. Each instance maintains its own routing table. The term Multiple Virtual Router (MVR) is sometimes used to represent a router that contains many VRF instances.

The Global Router, VRF 0, is the first instance of the router. When the system starts, it creates VRF 0 by default. VRF 0 provides all nonvirtual and traditional routing services. You cannot delete this instance. You can create and configure other VRF instances, if required.

VRF 0 is the only VRF that you can log into through ACLI. ACLI requires you to specify the VRF when you enter commands.

You can associate one VRF instance with many IP interfaces. These interfaces are unique for each VRF instance. An interface is an entity with an IP address that has the following characteristics:

- A unique association with a VLAN and VLAN ID
- A unique association with a brouter, if not associated with a VLAN
- A unique association with a circuit

A VLAN can only be associated with a single VRF instance.

Note:

You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IP address. You must first associate the port and VRF instance and then you can configure the IP address.

VRF Lite capability and functionality

The Avaya Virtual Services Platform 9000 supports what is termed VRF Lite. Lite conveys the fact that the device does not use Multiprotocol Label Switching (MPLS) for VRF; VRF Lite is a device virtualization feature, not a network-wide virtualization feature.

On a VRF instance, VRF Lite supports the following protocols:

- Border Gateway Protocol (BGP)
- IP
- Internet Control Message Protocol (ICMP)
- Address Resolution Protocol (ARP)
- Static routes
- Default routes
- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- Route policies
- Virtual Router Redundancy Protocol (VRRP)
- Dynamic Host Configuration Protocol (DHCP), and BootStrap Protocol relay agent
- User Datagram Protocol (UDP) forwarding
- Internet Group Management Protocol (IGMP)

All OSPF VRF instances use one OSPF timer, and all RIP VRF instances use one RIP timer.

Avaya Virtual Services Platform 9000 uses VRF Lite to perform the following actions:

- Partition traffic and data and represent an independent router in the network
- Provide virtual routers that are transparent to end-users
- Support addresses that are not restricted to the assigned address space provided by host Internet Service Providers (ISP)
- Support overlapping IP address spaces in separate VRF instances

Note:

If you enable multicast route redistribution between two VRFs, the switch does not support IP addresses that overlap within the two VRFs. The device does not generate an error if addresses overlap. You must avoid this situation.

VRF Lite interoperates with RFC 4364, Layer 3 VPNs. Split MultiLink Trunking (SMLT) and Routed SMLT (RSMLT) are also supported for VRF instances.

Although customer data separation into Layer 3 virtual routing domains is usually a requirement, sometimes customers must access a common network infrastructure. For example, they want to access the Internet, data storage, Voice over IP (VoIP)-public switched telephone network (PSTN), or call signaling services. To interconnect VRF instances, you can use an external firewall that supports virtualization, or use inter-VRF forwarding for specific services. With the inter-VRF solution, you can use routing policies and static routes to inject IP subnets from one VRF instance to another, and you can use filters to restrict access to certain protocols. The following figure depicts inter-VRF forwarding by Avaya Virtual Services Platform 9000.

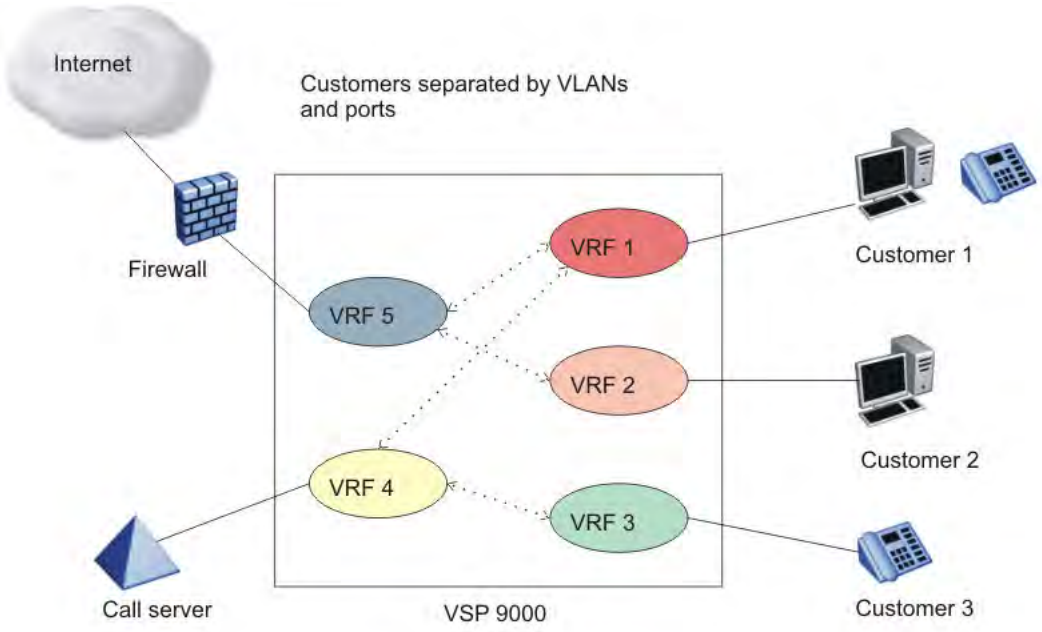


Figure 17: Inter-VRF forwarding

For more information about the latest VRF Lite scalability, see *Release Notes for Avaya Virtual Services Platform 9000,* NN46250-401.

For configuration information about multicast virtualization, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

VRF Lite and inter-VRF route redistribution

The Avaya Virtual Services Platform 9000 supports three route redistribution functions:

- intra-VRF inter-protocol route redistribution (redistribution within the same VRF instance), for example, redistribute RIP to OSPF
- inter-VRF inter-protocol redistribution (redistribution between two VRF instances), for example, redistribute RIP in VRF 2 to OSPF in VRF 4
- inter-VRF static routes (for example, a static route in a given VRF instance) configured as a typical static route but with the added parameter of a next-hop-vrf (the next-hop IP address is found in the next-hop-vrf instance)

With inter-VRF route redistribution, a user in one VRF instance can access route data in other VRF instances. You can redistribute routes within a VRF instance or between VRF instances; for example, one VRF instance can redistribute routes to all other VRF instances. You can redistribute Local, static, OSPF, RIP, and BGP routes and both dynamic (OSPF, BGP, and RIP) and static route redistribution is supported.

More than one routing protocol can be present in each VRF instance. Route redistribution can occur either between different protocol types, or between the same protocol types on different VRF instances.

An interface uses redistribution to announce routes that are learned by other protocols (OSPF or BGP, for example). Control route redistribution by using route policies. When you associate routing policies with route redistribution, the policy is checked before the target protocol is updated. Across VRF instances, the policy is checked at the source VRF instance, so only qualified routes are added to the routing table.

You can use static route commands to inject one specific route (including a default route) from one VRF instance to another. The route is added to the target VRF instance, while the next hop is resolved by the next-hop VRF instance.

Static routes are used to direct packets from a given source using a next-hop IP address. The next-hop-vrf option in a static route permits this path to proceed from one VRF to another. Overlapping IP addresses are supported within VRFs, thus it is possible for two VRFs to have identical IP addresses.

The following list describes interVRF route redistribution:

- Redistributed routes are added to the target VRF instance, and their next hop remains in the source VRF instance.
- If either the source or destination VRF instance is deleted, the redistribution configuration is automatically deleted.
- Redistributed routes are not further redistributed to another VRF instance.
- Route redistribution is unidirectional. You must configure route redistribution for the reverse direction if you require it. You can configure different route policies for each direction.

- After you configure interVRF route redistribution between two VRF instances, you must avoid using overlapping IP addresses between these two VRF instances.

  Avoid overlapping addresses; the device does not generate an error if addresses overlap.
- Intra-VRF routes take precedence over inter-VRF routes.
- You can physically connect two VRF instances to distribute route across VRF instances (in this case, you do not need to configure route redistribution).

Route redistribution operation

To perform redistribution, the device maintains a route change list. The change list contains all the best routes that are either added to or deleted from the forwarding table. When a best route is added to or deleted from the forwarding table, the change list is updated to reflect the change and notify registered protocols. The registered protocols pick up the change from the change list when it becomes available.

An example scenario of interVRF redistribution follows. To redistribute OSPF routes in VRF 1 to RIP in VRF 0:

- Create, enable, and apply a RIP redistribution instance. The source protocol is OSPF and the VRF source is VRF 1.
- When an OSPF route is added in VRF 1, the Routing Table Manager (RTM) in VRF 1 puts the new route into the change list.
- The device notifies RIP in VRF 0, because RIP is registered with the RTM of VRF 1 for OSPF route changes.
- To send OSPF routes from VRF 1 through the RIP interface in VRF 0, the interface uses a route policy with match VRF criterion of VRF 1.

Virtual Services Platform 9000 also supports inter-domain multicast routing. For more information, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

VRF Lite requirements

To use VRF Lite, you require the following hardware and software:

- Avaya Virtual Services Platform 9000 Software Release 3.0 or later
- Premier Software License

Ethernet modules and VRF Lite management

You can configure each VRF instance as a separate router, this means that you can configure different routing protocols and associated parameters for each instance. You can associate non0 VRF instances with module ports.

The Ethernet parameters that you can edit for a VRF instance depend on whether the port belongs to only one, or more than one, VRF instance. For example, if a port belongs to only one VRF, you

can edit the Ethernet parameters of the VRF. If a port belongs to more than one VRF instance, you cannot edit the Ethernet parameters of that port unless you are accessing the port through the Global Router with read-write-all access. If you do not have read-write-all access, you can only edit the GlobalRouter port parameters. If a port belongs to a single non0 VRF, the port Ethernet parameters can be changed by this VRF. If a port belongs to multiple VRF instances, only a user with read-write-all access who is accessing the port through the Global Router can change this port configuration.

Management VRF

Management port

The management port on the CP module is a 10/100/1000 Mb/s Ethernet port that you can use for an out-of-band management connection to the switch.

Configure IP addresses for management ports 1/1 and 2/1 to remotely access the switch using the management ports.

You can also configure a virtual management IP address for the management ports for redundancy, which gives you an alternate way to remotely connect to the management ports. The virtual IP address for the management ports is a floating address that is always owned by the current primary CPU. After a switchover, the IP address is automatically assigned to the new primary CPU, which allows you to use a single IP address to access the management interface regardless of which CPU is the primary.

The IP addresses for each of the management ports and the virtual management IP address must be on the same IP subnet.

Management Router VRF

Virtual Services Platform 9000 has a separate VRF called Management Router (MgmtRouter) reserved for OAM ports 1/1 and 2/1, and the Virtual Management IP address. The configured IP subnet has to be globally unique because the management protocols, for example, SNMP, Telnet, and FTP, can go through in-band or out-of-band ports. The VRF ID for the Management Router is 512.

Virtual Services Platform 9000 never switches or routes transit packets between the Management Router VRF port and the Global Router VRF, or between the Management Router VRF and other VRF ports.

Avaya honors the VRF of the ingress packet; however, in no circumstance does Virtual Services Platform 9000 allow routing between the Management VRF and Global Router VRF. VSP 9000 does not support the configuration if you have an out-of-band management network with access to the same networks present in the GRT routing table.

Non-virtualized client management applications

Avaya recommends that you do not define a default route in the Management Router VRF. A route used for non-virtualized client management applications, such as Telnet, Secure Shell (SSH), and FTP, originating from the VSP 9000, will always match a default route defined in the Management Router VRF.

If you want out-of-band management, Avaya recommends that you define a specific static route in the Management Router VRF to the IP subnet where your management application resides.

When you specify a static route in the Management Router VRF, it enables the client management applications originating from the VSP 9000 to perform out-of-band management without affecting in-band management. This enables in-band management applications to operate in the Global Router VRF.

Non-virtualized client management applications originating from the VSP 9000, such as Telnet, SSH, and FTP, follow the behavior listed below:

1. Look at the Management Router VRF route table
2. If no route is found, the applications will proceed to look in the Global Router VRF table

Non-virtualized client management applications include:

- DHCP Relay
- DNS
- FTP client with the **copy** command
- IPFIX
- NTP
- rlogin
- RADIUS authentication and accounting
- SSH
- SNMP clients in the form of traps
- SYSLOG
- TACACS+
- Telnet
- TFTP client

For management applications that originate outside the VSP 9000, the initial incoming packets establish a VRF context that limits the return path to the same VRF context.

Virtualized management applications

Virtualized management applications, such as ping and traceroute, operate using the specified VRF context. To operate ping or traceroute you must specify the desired VRF context. If not specified, ping defaults to the Global Router VRF. For example, if you want to ping a device through the out-of-band management port you must select the Management Router VRF.

```
VSP-9012:1(config)#ping 192.0.2.1 vrf MgmtRouter
192.0.2.1 is alive
```

VRF Lite configuration rules

You must select the VRF for global IP options before entering commands.

Not all Global Router parameters are configurable on other VRF instances.

For instructions about how to configure a VRF instance, see the following paragraphs. For instructions about how to configure other parameters, for example, OSPF, RIP, and ECMP for a VRF instance, see the instructions to configure a router for OSPF, RIP, and ECMP.

Layer 1 and Layer 2 information (including VLAN information) is global and is not maintained for each VRF instance. However, you can associate a set of VLANs with a VRF instance.

One VLAN cannot belong to more than one VRF instance at one time. When you create a VLAN, more than one physical port can belong to it. You can associate a VRF instance with more than one IP interface (a physical Ethernet port or a VLAN).

Perform physical port assignment at the VLAN and brouter port level. A VRF instance inherits all the ports assigned to its VLANs and brouter ports. You cannot directly assign a physical port to a VRF instance, but it is implicitly assigned when you associate the VRF with VLANs or brouter ports.

After you configure interVRF route redistribution between two VRF instances, avoid overlapping IP addresses between these two VRF instances.

To delete an OSPF instance, first disable OSPF, and then delete the OSPF instance.

To activate OSPF on an interface, first enable OSPF on the VRF instance (the VRF instance to which the interface is bound), and then enable OSPF on the interface. In non-0 VRF instances, RIP and OSPF must be created before they can be enabled.

When you configure VRF Lite, remember the following points:

- You cannot associate a brouter port or VLAN with a VRF instance if the brouter port or VLAN has an IP address. Configure the association first, and then configure required IP addresses.
- You cannot configure an IP interface (VLAN or brouter port) for a VRF instance until the VRF instance exists.
- You can delete a VRF instance only after you delete all its interfaces and other subcomponents.
- Before you delete a VRF instance, disable OSPF. Deleting a VRF instance deletes the OSPF instance if OSPF is disabled.
- When you create a VRF instance, an OSPF instance is not automatically created. To activate OSPF on a VRF instance, first create an OSPF instance, and then enable OSPF.
- An IP routable VLAN can become a member of a VRF.
- An IP interface can belong to only one VRF.
- A VRF can exist even if no interfaces are assigned to it.
- You can connect two VRFs from the same system with an external cable.
- Routing policies apply to VRFs on an individual basis.
- You can configure a VRF so it can have IP interfaces with OSPF, RIP, static routes, and policies simultaneously.
- VRF Lite supports RIP in and out policies.

- VRF Lite supports OSPF in and out (accept and redistribute) policies.
- Multiple VRFs on the same node can function in different autonomous systems.
- VRF Lite supports SMLT and RSMLT.
- If you configure an IP interface without specifying the VRF instance, it is mapped to VRF 0 by default.
- Every interface is a member of VRF 0 unless explicitly defined to belong to another VRF.

Virtualized protocols

VRF Lite supports virtualization of the following protocols and features. Use this table to find applicable VRF command and procedure information.

Table 63: Virtualized protocols and documentation

Virtualized protocol or feature	Where to find information
ARP	This document
BGP	*Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507
Circuitless IP	This document
DHCP	This document
IGMP	*Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504
OSPF	*Configuring OSPF and RIP on Avaya Virtual Services Platform 9000,* NN46250-506
RIP	*Configuring OSPF and RIP on Avaya Virtual Services Platform 9000,* NN46250-506
Route policies	This document
Route preferences	This document
Router Discovery	This document
Static routes	This document
User Datagram Protocol (UDP)	This document
VLAN	*Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500
VRRP	This document

Chapter 17: VRF Lite configuration using ACLI

Use Virtual Router and Forwarding (VRF) Lite to provide many virtual routers using one Virtual Services Platform 9000.

This section shows you how to configure a VRF instance and how to associate ports and VLANs with VRF instances. For instructions about how to configure other parameters, for example, Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and Equal Cost Multipath (ECMP) for a VRF, see the instructions to configure a router for OSPF, RIP, and ECMP.

- VRF Lite requires the Premier software license.

For more information about multicast virtualization, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

The following task flow shows you the sequence of procedures you perform to configure VRF Lite.

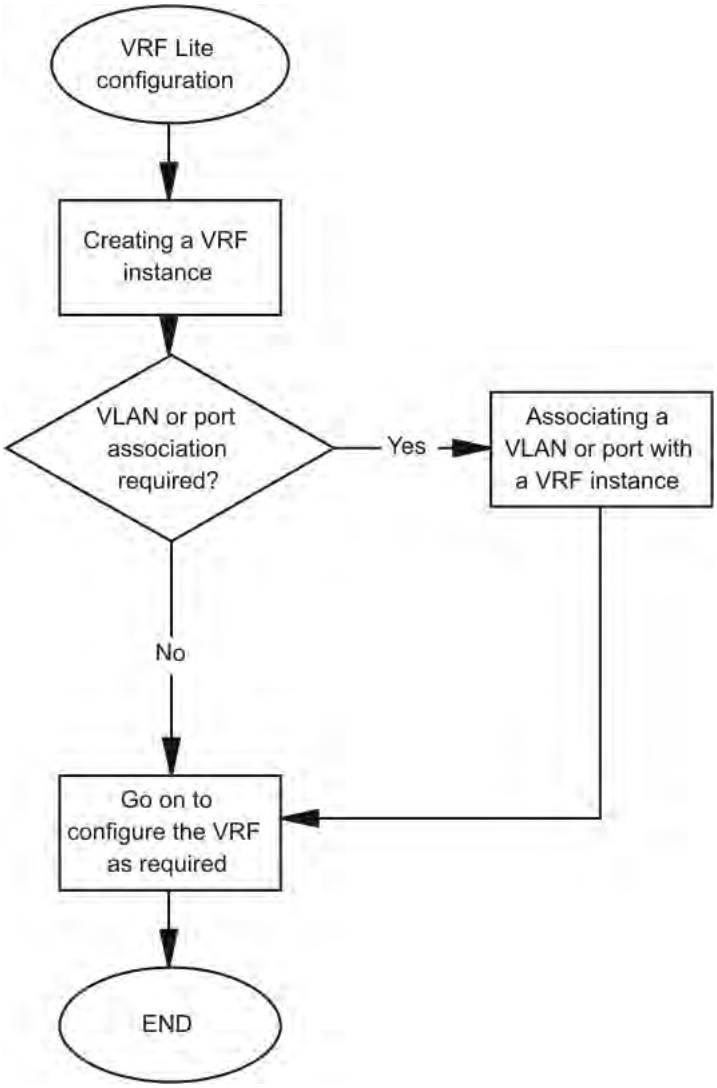


Figure 18: VRF Lite configuration procedures

Creating a VRF instance

Create a VRF instance to provide a virtual routing interface for a user.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create a VRF instance and specify a VRF name:

```
ip vrf WORD<1-16>
```

3. Configure the maximum number of routes:

   ```
   ip vrf WORD<1-16> max-routes <0-250000>
   ```

4. Enable max-routes traps:

   ```
   ip vrf WORD<1-16> max-routes-trap enable
   ```

5. Enter VRF Router Configuration mode for a specific VRF context:

   ```
   enable
   configure terminal
   router vrf WORD<1-16>
   ```

6. Configure the routing protocol triggers for the VRF:

   ```
   ip bgp
   ip ospf
   ip rip
   ip igmp
   ```

   You cannot configure BGP, OSPF, or RIP on a VRF instance unless you first configure the routing protocol trigger.

7. Ensure that the instance is configured correctly:

   ```
   show ip vrf [WORD<1-16>]
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Create a VRF instance and specify a VRF name:

```
VSP-9012:1(config)#ip vrf test1
```

Configure the maximum number of routes:

```
VSP-9012:1(config)#ip vrf test1 max-routes 12000
```

Enable max-routes traps:

```
VSP-9012:1(config)#router vrf test1 max-routes-trap enable
```

Enter Router Configuration mode:

```
VSP-9012:1(config)#router vrf test1
```

Configuration the routing protocol triggers for the VRF:

```
VSP-9012:1(router-vrf)#ip bgp
VSP-9012:1(router-vrf)#ip ospf
VSP-9012:1(router-vrf)#ip rip
```

```
VSP-9012:1(router-vrf)#ip igmp
```

Ensure that the instance is configured correctly:

```
VSP-9012:1(router-vrf)#show ip vrf test
```

Variable definitions

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

Table 64: Variable definitions

Variable	Value
max-routes <0-250000>	Specifies the maximum number of routes for the VRF. The default value is 10000, except for the Global Router, which is 250000.
max-routes-trap [enable]	Enables the sending of traps after the maximum number of routes is reached.
name <WORD 1-16>	Renames the VRF instance.
vrf-trap	Enables the device to send VRF-related traps.
vrfid<*1–511*>	Specifies the ID number for the VRF instance. VRF ID 0 is reserved for the GlobalRouter.
<WORD 1–16>	Specifies the VRF name.

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

Table 65: Variable definitions

Variable	Value
max-routes [vrfids WORD<1-512>] [<WORD 1-16>]	Displays the maximum number of routes for the specified VRFs. • vrfids WORD<0-512> specifies a list of VRFs by VRF IDs. • WORD<1-16> specifies a VRF by name.
vrfids WORD<1-512>	Specifies a list of VRFs by VRF IDs.
WORD<1-16>	Specifies a VRF by name.

Associating a VLAN or port with a VRF instance

You can assign a VRF instance to a port or VLAN. You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IP address. You can configure the IP address after you associate the port and VRF instance.

Before you begin

- Ensure the VRF is already configured.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   interface vlan <1-4084>
   ```

2. Associate the VLAN with a VRF instance:

   ```
   vrf WORD<1-16>
   ```

3. Log on to GigabitEthernet Interface Configuration mode.
4. Associate a port with a VRF instance:

   ```
   vrf <WORD 1-16>
   ```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
```

Create a VRF named Two:

```
VSP-9012:1(config-if)#ip vrf Two
```

Create a VLAN of type byport:

```
VSP-9012:1(config-if)#vlan create 33 name vlan-30 type port-mstprstp 0
```

Enter VLAN Interface Configuration mode:

```
VSP-9012:1(config-if)#interface vlan 33
```

Give the VLAN an IP address:

```
VSP-9012:1(config-if)#ip address 32.22.12.2 255.255.255.0
```

Assign the VLAN to VRF Two:

```
VSP-9012:1(config-if)#vrf Two
```

Enter VRF configuration mode:

```
VSP-9012:1(config-if)#router vrf Two
```

Enable OSPF on the VRF:

```
VSP-9012:1(config-if)#ip ospf
```

Variable definitions

Use the data in the following table to use the **vrf** command.

Table 66: Variable definitions

Variable	Value
vrf WORD<1-16>	Specifies the VRF instance by name.

Configuring a static route on the management router

Before you begin

- You must log on to the router VRF mgmtRouter mode in ACLI.

About this task

Configure a management static route. The management router only supports static routing.

Avaya recommends that you do not define a default route in the Management Router VRF. A route used for non-virtualized client management applications, such as Telnet, Secure Shell (SSH), and FTP, that originates from the VSP 9000, will always match a default route defined in the Management Router VRF.

Non-virtualized client management applications originating from the VSP 9000, such as Telnet, SSH, and FTP, first look at the Management Router VRF route table. If not route is found, the applications will proceed to look in the Global Router VRF table.

If you want in-band management, Avaya recommends that you define a specific static route in the Management Router VRF to the IP subnet where your management application resides. When you specify a static route in the Management Router VRF, it enables the client management applications originating from the VSP 9000 to perform out-of-band management without affecting in-band management. This enables in-band management applications to operate in the Global Router VRF.

The NH VRF/ISID column in the **`show ip route [vrf WORD<1-16>]`** command displays the I-SID in the following examples:

- Only for inter-Virtual Services Network (VSN) routes leaked using IS-IS accept policies.
- Only if the I-SID for which the routes are leaked does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays.

If the I-SID is 0, which represents the GlobalRouter, the column displays as GlobalRouter. The existing IS-IS routes in Shortest Path Bridging (SPB) Layer 3 VSN continue to display as the VRF name of the IP VSN.

Procedure

1. Configure the static route:

   ```
   ip route {A.B.C.D} {A.B.C.D} {A.B.C.D} weight <1-65535>
   ```

2. Confirm that the static route is correct:

   ```
   show ip route vrf mgmtrouter
   ```

Example

```
VSP-9012:1>enable
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Log on to the router VRF mgmtRouter mode:

```
VSP-9012:1(config)#router vrf mgmtRouter
```

Configure the static route:

```
VSP-9012:1(config-vrf)#ip route 46.0.0.0 255.0.0.0 46.17.159.128 weight 1
```

Enable the static route:

```
VSP-9012:1(config-vrf)#ip route 46.0.0.0 255.0.0.0 46.17.159.128 enable
```

Confirm the static route is correct:

```
VSP-9012:1(config-vrf)#show ip route vrf mgmtrouter
```

```
VSP-9012:1(config)#show ip route vrf mgmtrouter

================================================================================
                             IP Route - VRF MgmtRouter
================================================================================

                                       NH             INTER
DST           MASK          NEXT       VRF/ISID COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------

192.0.2.0    255.0.0.0     192.0.2.128 MgmtRouter 1  4092 STAT 0   IB   5
198.51.100.1 255.255.254.0 192.0.2.91  -          1  4092 LOC  0   DB   0

2 out of 2 Total Num of Route Entries, 2 Total Num of Dest Networks displayed.
--------------------------------------------------------------------------------

TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp
Route,U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route,
S=SPBM Route
```

Variable definitions

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

Variable	Value
{A.B.C.D} {A.B.C.D} {A.B.C.D}	Specifies the IP address, subnet mask, and the next-hop address for the route.
weight <*1–65535*>	Specifies the cost of hops for the route.

Creating an IP VPN instance on a VRF

Create an IP VPN instance to advertise IP routes from one VRF to another across a Shortest Path Bridging MAC (SPBM) network.

For more information about Layer 3 Virtual Services Networks (VSN) and SPBM, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Before you begin

- The VRF must exist.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

   ```
   enable
   configure terminal
   router vrf WORD<1-16>
   ```

2. Create an IP VPN instance on the VRF:

   ```
   ipvpn
   ```

3. Assign a service instance identifier (I-SID) to the IP VPN:

   ```
   i-sid <0-16777215>
   ```

4. Enable IP VPN on the VRF:

   ```
   ipvpn enable
   ```

   By default, a new IP VPN instance is disabled.

5. Display all IP VPNs:

   ```
   show ip ipvpn [vrf WORD<1-16>] [vrfids WORD<0-512>]
   ```

Example

From Global Configuration mode, log on to Router VRF Configuration mode:

```
VSP-9012:1(config)#router vrf red
```

Create the IP VPN instance:

```
VSP-9012:1(router-vrf)#ipvpn
```

Enable IP VPN:

```
VSP-9012:1(router-vrf)#i-sid 100
```

Enable IP VPN:

```
VSP-9012:1(router-vrf)#ipvpn enable
```

```
VSP-9012:1#show ip ipvpn
          VRF Name             : red
          Ipvpn-state          : enabled
          I-sid                : 100
```

Variable definitions

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

Variable	Value
vrf *WORD<1–16>*	Specifies the VRF name.
vrfids *WORD<0–512>*	Specifies the VRF ID.

Use the data in the following table to use the **`i-sid`** command.

Variable	Value
i-sid *<0–16777215>*	Assigns an I-SID to the VRF to configure. Use the no or default option to remove the I-SID to VRF allocation for this VRF.

Chapter 18: VRF Lite configuration using Enterprise Device Manager

Use VRF Lite to provide many virtual routers using one Avaya Virtual Services Platform 9000.

For more information about multicast virtualization, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

Configuring a VRF instance

About this task

Configure a VRF instance to provide a virtual routing interface for a user.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRF**.
3. Click **Insert**.
4. Specify the VRF ID.
5. Name the VRF instance.
6. To enable the VRF to send VRF Lite-related traps, select **TrapEnable**.
7. Configure the other parameters as required.
8. Click **Insert**.

VRF field descriptions

Use the data in the following table to help you use the **VRF** tab.

Name	Description
Id	Specifies the ID number of the VRF instance. VRF ID 0 is reserved for the GlobalRouter.
Name	Names the VRF instance.

Name	Description
ContextName	Identifies the VRF. The SNMPv2 Community String or SNMPv3 contextName denotes the VRF context and is used to logically separate the MIB module management.
TrapEnable	Enables the VRF to send VRF Lite-related traps (VrfUp and VrfDown). The default is enabled.
MaxRoutes	Configures the maximum number of routes allowed for the VRF. The default value is 10000, except for the GlobalRouter, which is 250000.
RpTrigger	Specifies the Routing Protocol (RP) triggers for the VRF. The triggers are used to initiate or shutdown routing protocols on a VRF. The protocols include RIP, OSPF and BGP. You can act upon multiple RPs simultaneously. Also, you can use this option to bring individual RPs up in steps.
MaxRoutesTrapEnable	Enables the generation of the VRF Max Routes Exceeded traps. The default is enabled.

Configuring interVRF route redistribution policies

Before you begin

- Ensure VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required.

About this task

Configure inter-VRF route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF or BGP. Use a route policy to control the redistribution of routes.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **Policy**.
3. Click the **Route Redistribution** tab.
4. Click **Insert**.
5. Choose the source and destination VRF IDs.
6. Choose the protocol and route source.
7. Select **Enable**.
8. Choose the route policy to apply to the redistributed routes.
9. Configure other parameters as required.
10. Click **Insert**.

11. Click the **Applying Policy** tab.
12. Select **RedistributeApply**, and then click **Apply**.

Route Redistribution field descriptions

Use the data in the following table to use the **Route Redistribution** tab.

Name	Description
DstVrfId	Specifies the destination VRF ID to use in the redistribution.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrfId	Specifies the source VRF ID to use in the redistribution.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution. The default is disabled.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine whether the system advertises a specific route to the specified protocol.
Metric	Specifies the metric announced in advertisements. The default is 0.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone. The default is type2.
Subnets	Indicates that subnets must be advertised individually (applies to OSPF only). The default is allow.

Viewing brouter port and VRF associations

About this task

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

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRF**.
3. Click the **VRF-Ports** tab.
4. To display the VRF names associated with a port , click a cell in one of the table rows and, on the toolbar, click the **ShowVRFNames** button.

5. To change the VRF, double-click the **BrouterVrfId** field for the port.

   Tip:

   You can associate a port with more than one VRF.

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

VRF-Ports field descriptions

Use the data in the following table to use the **VRF-Ports** tab.

Name	Description
Index	Specifies the slot and port.
Type	Specifies the port type.
VrfIds	Identifies the set of VRF IDs to which this port belongs.
VrfCount	Shows the number of VRF instances associated with this port.
BrouterVrfId	Shows the VRF ID for this brouter port.
BrouterVrfName	Shows the VRF name for this brouter port.
Show VrfNames	You can use this toolbar button to identify the set of VRF names to which a port belongs.

Use the data in the following table to use the **Show VrfNames** button.

Name	Description
Index	Specifies the slot and port.
VrfNames	Shows the VRF name for this brouter port.

Viewing global VRF status information

About this task

View global VRF status information to determine the number of VRFs that are configured and active.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRF**.
3. Click the **Global Status** tab.

Global Status field descriptions

Use the data in the following table to use the **Global Status** tab.

Name	Description
ConfigNextAvailableVrfId	Specifies the number of the next available Virtual Router ID (index).
ConfiguredVRFs	Specifies the number of VRFs configured on this network element.
ActiveVRFs	Specifies the number of VRFs that are active on the network element. These are VRFs for which the OperStatus is up.

Viewing VRF instance statistics and status information

About this task

View VRF instance status information to determine the operational status of each VRF, as well as other operational parameters.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP** .
2. Click **VRF**.
3. Click the **VRF Stats** tab.

VRF Stats field descriptions

Use the data in the following table to use the **VRF Stats** tab.

Name	Description
Id	Specifies the ID number of the VRF instance.
StatRouteEntries	Specifies the total number of routes for this VRF.
StatFIBEntries	Specifies the total number of Forwarding Information Base (FIB) entries for this VRF.
StatUpTime	Specifies the time in (in hundredths of a second) since this VRF entry has been operational.
OperStatus	Shows the operational status of the Virtual Router.
RpStatus	Shows the status of the routing protocols used on this VRF that correspond to the list specified in VrfRpTrigger.
RouterAddressType	Specifies the router address type of this VRF.

Name	Description
Router Address	Specifies the router address of this VRF, derived from one of the interfaces. If a loopback interface is present, you can use the loopback interface address.

Viewing VRF statistics for a VRF

About this task

View VRF statistics to ensure the instance is performing as expected.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **VRF**.
3. Click the **VRF** tab.
4. Select a VRF.
5. Click the **VRF Stats** button.

Stats field descriptions

Use the data in the following table to help you understand the VRF statistics.

Name	Description
StatRouteEntries	Specifies the total number of routes for this VRF.
FIBEntries	Specifies the total number of Forwarding Information Base (FIB) entries for this VRF.

Selecting and launching a VRF context view

Use this procedure to switch to a VRF context view and launch it so that you can view and configure features for the VRF instance.

About this task

Global Router 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

1. In the navigation pane, expand the following folders: **Configuration** > **VRF Context View**.
2. Click **Set VRF Context View**.
3. Click the **VRF** tab.
4. Select a context to view.
5. Click **Launch VRF Context view**.

   A new browser tab appears that contains the selected VRF view.

VRF field descriptions

Use the data in the following table to use the **VRF** tab.

Name	Description
Id	Shows the unique VRF ID.
Name	Shows the name of the virtual router.
ContextName	Shows the SNMPv3 context name that denotes the VRF context and logically separates the MIB module management.

Creating an IP VPN instance on a VRF

Create an IP VPN instance to advertise IP routes from one VRF to another across a Shortest Path Bridging MAC (SPBM) network.

For more information about Layer 3 Virtual Services Networks (VSNs) and SPBM, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- The VRF must exist.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
2. Click **IP-VPN**.
3. Click the **VPN** tab.
4. Click **Insert**.
5. Click the ellipsis button **[...]**, and then select a VRF from the list.
6. Click **OK**.
7. Click **Insert**.

   By default, the new IP VPN instance is disabled.
8. In the **IsidNumber** column, double-click the **0** value, and then enter the service instance identifier (I-SID) to assign to the IP-VPN.
9. In the **Enable** column, double-click the **disable** value.
10. Click the arrow to view a list of choices, and then choose **enable**.
11. Click **Apply**.

VPN field descriptions

Use the data in the following table to use the **VPN** tab.

Name	Description
VrfId	Specifies the ID of the VRF to configure.
Enable	Enables or disables the IP VPN instance on the VRF. The default is disabled.
IsidNumber	Specifies the I-SID to associate with the VPN. By default, no I-SID is assigned.

Glossary

Address Resolution Protocol (ARP)	Maps an IP address to a physical machine address, for example, maps an IP address to an Ethernet media access control (MAC) address.
aggregate	A prefix length that is formed by combining several specific prefixes. The resulting prefix is used to combine blocks of address space into a single routing announcement.
Autonomous System (AS)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the AS, and using an EGP to route packets to other ASs.
Autonomous System Number (ASN)	A two-byte number that is used to identify a specific AS.
Bootstrap Protocol (BootP)	A User Datagram Protocol (UDP)/Internet Protocol (IP)-based protocol that a booting host uses to configure itself dynamically and without user supervision.
bootstrap router (BSR)	A dynamically elected Protocol Independent Multicast (PIM) router that collects information about potential Rendezvous Point routers and distributes the information to all PIM routers in the domain.
Bridge Protocol Data Unit (BPDU)	A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.
candidate bootstrap router (C-BSR)	Provides backup protection in case the primary rendezvous point (RP) or boostrap router (BSR) fails. Protocol Independent Multicast (PIM) uses the BSR and C-BSR.
Circuitless IP (CLIP)	A CLIP is often called a loopback and is a virtual interface that does not map to any physical interface.
classless interdomain routing (CIDR)	The protocol defined in RFCs 1517 and 1518 for using subnetwork masks, other than the defaults for IP address classes.
Control Processor Unit High Availability (CPU-HA)	CPU-HA activates two CP modules simultaneously. The CP modules exchange topology data so, if a failure occurs, either CP module can take precedence in less than one second with the most recent topology data.

Dynamic Random Access Memory (DRAM)	A read-write random-access memory, in which the digital information is represented by charges stored on the capacitors and must be repeatedly replenished to retain the information.
Enterprise Device Manager (EDM)	A Web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.
equal cost multipath (ECMP)	Distributes routing traffic among multiple equal-cost routes.
Global routing engine (GRE)	The base router or routing instance 0 in the Virtual Routing and Forwarding (VRF).
Institute of Electrical and Electronics Engineers (IEEE)	An international professional society that issues standards and is a member of the American National Standards Institute, the International Standards Institute, and the International Standards Organization.
Interior Gateway Protocol (IGP)	Distributes routing information between routers that belong to a single Autonomous System (AS).
Internet Assigned Numbers Authority (IANA)	The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.
Internet Control Message Protocol (ICMP)	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.
Internet Protocol version 4 (IPv4)	The protocol used to format packets for the Internet and many enterprise networks. IPv4 provides packet routing and reassembly.
interswitch trunking (IST)	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.
Layer 1	Layer 1 is the Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interacts with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).

Layer 3 Virtual Services Network	The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN).
link-state advertisement (LSA)	Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
mask	A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.
maximum transmission unit (MTU)	The largest number of bytes in a packet—the maximum transmission unit of the port.
media	A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.
Media Access Control (MAC)	MAC arbitrates access to and from a shared medium.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.
multiplexing	Carriage of multiple channels over a single transmission medium; a process where a dedicated circuit is shared by multiple users. Typically, data streams intersperse on a bit or byte basis (time division), or separate by different carrier frequencies (frequency division).
Network Basic Input/ Output System (NetBIOS)	An application programming interface (API) that augments the DOS BIOS by adding special functions for Local Area Networks (LAN).
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
not so stubby area (NSSA)	Prevents the flooding of external link-state advertisements (LSA) into the area by providing them with a default route. An NSSA is a configuration of the Open Shortest Path First (OSPF) protocol.

operation, administration, and maintenance (OA&M)	All the tasks necessary for providing, maintaining, or modifying switching system services.
Packet Capture Tool (PCAP)	A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.
port	A physical interface that transmits and receives data.
prefix	A group of contiguous bits, from 0 to 32 bits in length, that defines a set of addresses.
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
Protocol Independent Multicast, Sparse Mode (PIM-SM)	PIM-SM is a multicast routing protocol for IP networks. PIM-SM provides multicast routing for multicast groups that can span wide-area and inter-domain networks, where receivers are not densely populated. PIM-SM sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.
remote monitoring (RMON)	A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.
Reverse Address Resolution Protocol (RARP)	A protocol that maintains a database of mappings between physical hardware addresses and IP addresses.
reverse path checking (RPC)	Prevents packet forwarding for incoming IP packets with incorrect or forged (spoofed) IP addresses.
route flapping	An instability that is associated with a prefix, where the associated prefix routes can exhibit frequent changes in availability over a period of time.
route table manager (RTM)	Determines the best route to a destination based on reachability, route preference, and cost.
Routed Split MultiLink Trunking (RSMLT)	Provides full router redundancy and rapid failover in routed core SMLT networks and as RSMLT-edge in routed SMLT edge applications; eliminating routing protocol timer dependencies when network failures occur.

Routing Information Protocol (RIP)	A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.
routing policy	A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.
Service Instance Identifier (I-SID)	The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 24 bits. SPBM uses this I-SID to identify and transmit any virtualized traffic in an encapsulated SPBM frame. SPBM uses I-SIDs to virtualize VLANs (Layer 2 Virtual Services Network [VSN]) or VRFs (Layer 3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone. With Layer 2 VSNs, you associate the I-SID with a customer VLAN, which is then virtualized across the backbone. With Layer 3 VSNs, you associate the I-SID with a customer VRF, which is also virtualized across the backbone.
Shortest Path Bridging (SPB)	Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.
Shortest Path Bridging MAC (SPBM)	Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.
Simple Network Management Protocol (SNMP)	SNMP administratively monitors network performance through agents and management stations.
SMLT aggregation switch	One of two IST peer switches that form a split link aggregation group. It connects to multiple wiring closet switches, edge switches, or customer premise equipment (CPE) devices.
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states

	that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Spanning Tree Group (STG)	A collection of ports in one spanning-tree instance.
time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
Trivial File Transfer Protocol (TFTP)	A protocol that governs transferring files between nodes without protection against packet loss.
trunk	A logical group of ports that behaves like a single large port.
Universal/Local (U/L)	Determines global and local link addresses; used with the Extended Unique Identifier (EUI).
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
variable-length subnet masking (VLSM)	Allocating IP addressing resources to subnets according to their individual need rather than some general network-wide rule.
virtual router	An abstract object managed by the Virtual Router Redundancy Protocol (VRRP) that acts as a default router for hosts on a shared LAN.
virtual router forwarding (VRF)	Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.
Voice over IP (VOIP)	The technology that delivers voice information in digital form in discrete packets using the Internet Protocol (IP) rather than the traditional circuit-committed protocols of the public switched telephone network (PSTN).