



Avaya Virtual Services Platform 4000 Configuration — IP Routing

Release 5.1
NN46251-505
Issue 08.03
September 2016

© 2016, 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: <https://support.avaya.com/helpcenter/getGenericDetails?detailId=C20091120112456651010> under the link "Warranty & Product Lifecycle" 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.

"Hosted Service" means a hosted service subscription that You acquire from either Avaya or an authorized Avaya Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Avaya or Avaya Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES IF YOU PURCHASE A HOSTED SERVICE SUBSCRIPTION FROM AVAYA OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE AVAYA WEBSITE, [HTTPS://SUPPORT.AVAYA.COM/LICENSEINFO](https://support.avaya.com/LicenseInfo) UNDER THE LINK "Avaya Terms of Use for Hosted Services" OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE, BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE

TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE. YOUR USE OF THE HOSTED SERVICE SHALL BE LIMITED BY THE NUMBER AND TYPE OF LICENSES PURCHASED UNDER YOUR CONTRACT FOR THE HOSTED SERVICE, PROVIDED, HOWEVER, THAT FOR CERTAIN HOSTED SERVICES IF APPLICABLE, YOU MAY HAVE THE OPPORTUNITY TO USE FLEX LICENSES, WHICH WILL BE INVOICED ACCORDING TO ACTUAL USAGE ABOVE THE CONTRACT LICENSE LEVEL. CONTACT AVAYA OR AVAYA'S CHANNEL PARTNER FOR MORE INFORMATION ABOUT THE LICENSES FOR THE APPLICABLE HOSTED SERVICE, THE AVAILABILITY OF ANY FLEX LICENSES (IF APPLICABLE), PRICING AND BILLING INFORMATION, AND OTHER IMPORTANT INFORMATION REGARDING THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTPS://SUPPORT.AVAYA.COM/LICENSEINFO](https://support.avaya.com/LicenseInfo), UNDER THE LINK "AVAYA SOFTWARE LICENSE TERMS (Avaya Products)" 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.

Licence 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 is the software contained within the list of Heritage Nortel Products located at <https://support.avaya.com/LicenseInfo> under the link "Heritage

Nortel Products” or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants Customer a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or 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.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

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: <https://support.avaya.com/Copyright> or such successor site as designated by Avaya. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE <HTTP://WWW.MPEGLA.COM>.

Service Provider

THE FOLLOWING APPLIES TO AVAYA CHANNEL PARTNER’S HOSTING OF AVAYA PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN AVAYA CHANNEL PARTNER’S HOSTING OF AVAYA PRODUCTS MUST BE AUTHORIZED IN WRITING BY AVAYA AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT

SOFTWARE OR CODECS, THE AVAYA CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS, AT THE AVAYA CHANNEL PARTNER’S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER.

WITH RESPECT TO CODECS, IF THE AVAYA CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE AVAYA CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE AVAYA CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE <WWW.SIPRO.COM/CONTACT.HTML>. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE <HTTP://WWW.MPEGLA.COM>.

Compliance with Laws

Customer acknowledges and agrees that it is responsible for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Avaya product is used.

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: <https://support.avaya.com> or such successor site as designated by Avaya.

Security Vulnerabilities

Information about Avaya’s security support policies can be found in the Security Policies and Support section of <https://support.avaya.com/security>.

Suspected Avaya product security vulnerabilities are handled per the Avaya Product Security Support Flow (<https://support.avaya.com/css/P8/documents/100161515>).

Downloading Documentation

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

Contact Avaya Support

See the Avaya Support website: <https://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: <https://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	10
Purpose.....	10
Related resources.....	11
Support.....	13
Chapter 2: New in this document	15
Chapter 3: IP routing operations fundamentals	16
IP addressing.....	16
Loopback.....	20
IPv6 Circuitless IP.....	21
Static routes.....	21
Black hole static routes.....	22
VLANs and routing.....	23
Equal Cost Multipath.....	24
Alternate route.....	24
Route filtering and IP policies.....	26
Prefix list.....	29
Route policy definition.....	30
Address Resolution Protocol.....	34
Reverse Address Resolution Protocol.....	36
DHCP option 82.....	37
UDP broadcast forwarding.....	39
Virtual Router Redundancy Protocol.....	40
VRRP guidelines.....	43
VRRPv3.....	47
VRRPv3 guidelines.....	47
Layer 3 switch clustering and multicast SMLT.....	48
General guidelines.....	49
Multicast triangle topology.....	51
Square and full-mesh topology multicast guidelines.....	52
SMLT and multicast traffic issues.....	55
RSMLT.....	57
Enable or disable IPv4 ICMP broadcast.....	60
Chapter 4: ARP configuration using ACLI	61
Enabling ARP on a port or a VLAN.....	61
Enabling ARP proxy.....	62
Showing ARP information.....	62
Configuring IP ARP static entries.....	64
Clearing ARP entries.....	66
Showing ARP table information.....	67

Configuring Gratuitous ARP.....	69
Chapter 5: ARP configuration using Enterprise Device Manager.....	71
Enabling or disabling ARP on the brouter port or a VRF instance.....	71
Enabling or disabling ARP on a VLAN or a VRF instance.....	72
Viewing and managing ARP.....	73
Creating static ARP entries.....	74
Configuring ARP proxy.....	74
Chapter 6: DHCP and UDP configuration using ACLI.....	76
Configuring DHCP parameters globally.....	76
Showing DHCP relay information.....	78
Configuring DHCP option 82.....	79
Configuring DHCP relay on a port or VLAN.....	81
Configuring UDP broadcast forwarding.....	83
Configuring UDP protocols.....	83
Configuring a UDP port forward entry.....	84
Configuring the UDP port forwarding list.....	85
Showing UDP forward information.....	88
Chapter 7: DHCP and UDP configuration using Enterprise Device Manager.....	90
Configuring DHCP on a brouter port or a VRF instance.....	90
Configuring BootP/DHCP on a VLAN or VRF instance.....	92
Configuring DHCP relay.....	93
Viewing DHCP relay configuration information.....	94
Managing UDP forwarding protocols.....	95
Managing UDP forwarding.....	96
Creating the forwarding profile.....	97
Managing the broadcast interface.....	98
Viewing UDP endpoint information.....	99
Chapter 8: IP policy configuration using ACLI.....	101
Configuring prefix lists.....	101
Configuring an IPv6 prefix list.....	103
Configuring IP route policies.....	104
Configuring a policy to accept external routes from a router.....	111
Applying OSPF accept policy changes.....	113
Configuring inter-VRF redistribution policies.....	114
Chapter 9: IP policy configuration using Enterprise Device Manager.....	117
Configuring a prefix list.....	117
Configuring IPv6 Prefix List.....	118
Configuring a route policy.....	119
Applying a route policy.....	123
Viewing IP routes.....	124
Configuring an OSPF accept policy.....	126
Configuring inbound/outbound filtering policies on a RIP interface.....	127
Deleting inbound/outbound filtering policies on a RIP interface.....	128

Chapter 10: IP routing configuration using ACLI	129
Enabling routing globally or on a VRF instance.....	129
Enabling routing on an IP interface.....	131
Deleting a dynamically learned route.....	131
Configuring IP route preferences.....	132
Flushing routing tables by VLAN or port.....	134
Assigning an IP address to a port.....	134
Assigning an IP address to a VLAN.....	136
Viewing IP addresses for all router interfaces.....	137
Configuring IP routing globally or for a VRF.....	138
Configuring static routes.....	141
Configuring a black hole static route.....	144
Configuring a default static route.....	145
Enabling ICMP Router Discovery globally.....	147
Enabling or disabling IPv4 ICMP broadcast globally.....	148
Enabling or disabling IPv4 ICMP broadcast per VRF.....	149
Configuring Router Discovery on a port or VLAN.....	149
Configuring a CLIP interface.....	151
Creating an IPv6 CLIP interface.....	153
Variable definitions.....	154
Chapter 11: IP routing configuration using Enterprise Device Manager	155
Enabling routing for a router or a VRF instance.....	155
Deleting a dynamically-learned route.....	155
Configuring IP route preferences.....	157
Flushing routing tables by VLAN.....	158
Flushing routing tables by port.....	159
Assigning an IP address to a port.....	159
Assigning an IP address to a VLAN.....	161
Viewing IP addresses for all router interfaces.....	161
Configuring IP routing features globally.....	162
Configuring ECMP globally.....	165
Enabling alternative routes globally.....	166
Configuring static routes.....	166
Deleting a static route.....	168
Configuring a default static route.....	168
Configuring a black hole static route.....	169
Configuring ICMP Router Discovery globally.....	170
Configuring the ICMP Router Discovery table.....	170
Configuring ICMP Router Discovery for a port.....	172
Configuring ICMP Router Discovery on a VLAN.....	173
Configuring a Circuitless IPv4 interface.....	174
Enabling OSPF on a CLIP interface.....	175
Viewing TCP global information.....	176

Viewing TCP connections information.....	177
Viewing TCP listeners information.....	178
Chapter 12: RSMLT configuration using ACLI.....	180
Configuring RSMLT on a VLAN.....	180
Showing IP RSMLT information.....	181
Configuring RSMLT edge support.....	184
Chapter 13: RSMLT configuration using Enterprise Device Manager.....	186
Configuring RSMLT on a VLAN.....	186
Viewing and editing RSMLT local information.....	187
Viewing RSMLT peer information.....	188
Enabling RSMLT Edge support.....	189
Viewing RSMLT edge support information.....	190
Chapter 14: VRRP configuration using ACLI.....	191
Configuring VRRP on a port or a VLAN.....	191
Showing VRRP information	195
Showing extended VLAN VRRP.....	198
Showing VRRP interface information	199
Enabling ping to a virtual IP address.....	202
Configuring VRRP notification control.....	203
Configuring VRRP version on an interface.....	204
Enabling IPv4 VRRP preempt-mode.....	205
Chapter 15: VRRP configuration using EDM.....	207
Enabling VRRP global variables.....	208
Modifying VRRP parameters for an interface.....	208
Configuring VRRP on a V3 interface.....	211
Configuring VRRPv3 Checksum.....	213
Configuring Fast Advertisement Interval on a port or a VRF instance.....	214
Configuring Fast Advertisement Interval on a VLAN or a VRF instance.....	215
Chapter 16: VRF Lite fundamentals.....	216
Overview.....	216
VRF Lite capability and functionality.....	217
VRF Lite and inter-VRF route redistribution.....	218
VRF Lite requirements.....	220
Port parameters and VRF Lite management.....	220
VRF Lite configuration rules.....	220
Virtualized protocols.....	221
VRF Lite architecture examples.....	222
Chapter 17: VRF Lite configuration using ACLI.....	225
Creating a VRF instance.....	226
Associating a VLAN or port with a VRF instance.....	228
Creating an IP VPN instance on a VRF.....	230
Chapter 18: VRF Lite configuration using Enterprise Device Manager.....	232

Configuring a VRF instance.....	232
Configuring interVRF route redistribution policies.....	233
Viewing brouter port and VRF associations.....	234
Viewing global VRF status information.....	235
Viewing VRF instance statistics and status information.....	236
Viewing VRF statistics for a VRF.....	237
Selecting and launching a VRF context view	237
Creating an IP VPN instance on a VRF.....	238

Chapter 1: Introduction

Purpose

This document provides information on features in VSP Operating System Software (VOSS). VOSS runs on the following product families:

- Avaya Virtual Services Platform 4000 Series
- Avaya Virtual Services Platform 7200 Series
- Avaya Virtual Services Platform 8000 Series

This document provides procedures and conceptual information that you can use to configure the general routing operations for Avaya Virtual Services Platform 4000 Series. The operations included are:

- Address Resolution Protocol (ARP)
- Reverse ARP
- TCP and UDP
- Dynamic Host Configuration Protocol (DHCP) Relay
- Virtual Router Redundancy Protocol (VRRP)
- VRF-Lite
- Routed Split Multi-Link Trunking (RSMLT)
- Circuitless IP (CLIP) interfaces
- Static routes
- Point-to-Point Protocol over Ethernet
- Equal Cost Multipath (ECMP)
- Routing policies

For information to configure general routing operations on Avaya Virtual Services Platform 7200 Series and 8000 Series switches, see *Configuring IP Routing on Avaya Virtual Services Platform 7200 Series and 8000 Series*, NN47227-505.

Related resources

Documentation

See the *Documentation Roadmap for Avaya Virtual Services Platform 4000 Series*, NN46251-100 for a list of the documentation for this product.

Training

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

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.

Subscribing to e-notifications

Subscribe to e-notifications to receive an email notification when documents are added to or changed on the Avaya Support website.

About this task

You can subscribe to different types of general notifications, for example, Product Correction Notices (PCN), which apply to any product or a specific product. You can also subscribe to specific types of documentation for a specific product, for example, Application & Technical Notes for Virtual Services Platform 7000.

Procedure


1. In an Internet browser, go to <https://support.avaya.com>.
2. Type your username and password, and then click **Login**.
3. Under **My Information**, select **SSO login Profile**.
4. Click **E-NOTIFICATIONS**.
5. In the GENERAL NOTIFICATIONS area, select the required documentation types, and then click **UPDATE**.

The screenshot shows a web interface for managing notifications. At the top, it says 'GENERAL NOTIFICATIONS' and '1/5 Notifications Selected'. Below this is a list of notification types, each with a checkbox:

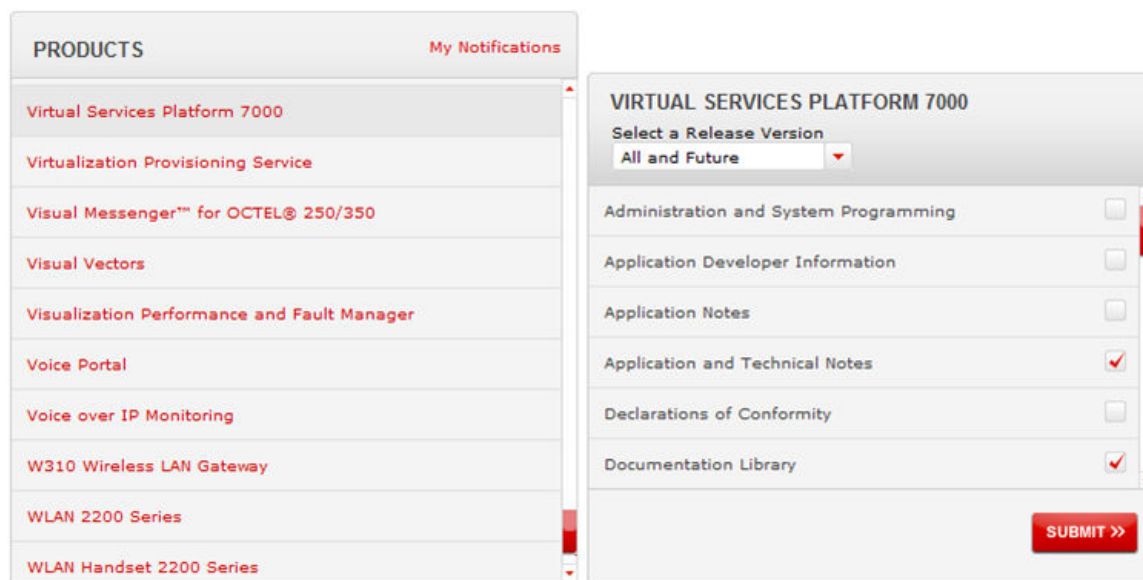
Notification Type	Selected
End of Sale and/or Manufacturer Support Notices	<input type="checkbox"/>
Product Correction Notices (PCN)	<input checked="" type="checkbox"/>
Product Support Notices	<input type="checkbox"/>
Security Advisories	<input type="checkbox"/>
Services Support Notices	<input type="checkbox"/>

At the bottom right of the form is a red button labeled 'UPDATE >>'.

6. Click **OK**.
7. In the PRODUCT NOTIFICATIONS area, click **Add More Products**.



8. Scroll through the list, and then select the product name.
9. Select a release version.
10. Select the check box next to the required documentation types.



11. Click **Submit**.

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 document

The following sections detail what is new in VOSS 5.1 in *Configuration - IP Routing for Avaya Virtual Services Platform 4000 Series*, NN46251-505.

Enable or disable IPv4 ICMP broadcast

This release introduces the option to enable or disable IPv4 ICMP broadcast processing.

For more information, see:

- [Enable or disable IPv4 ICMP broadcast](#) on page 60
- [Enabling or disabling IPv4 ICMP broadcast globally](#) on page 148
- [Enabling or disabling IPv4 ICMP broadcast per VRF](#) on page 149
- [Configuring IP routing features globally](#) on page 162

VRRPv3

VRRPv3 is a combined protocol for both IPv4 and IPv6. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IPv4 or IPv6 addresses associated with a virtual router is called Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Backups wait for a Master and take ownership when the Master is no longer detected.

The following VRRPv3 enhancements are available as part of this release:

1. Adding VRRPv3 for IPv4.
2. Making both IPv4 and IPv6 VRRPv3 features compliant to RFC5798.

For more information see:

- [VRRPv3](#) on page 47
- [Enabling IPv4 VRRP preempt-mode](#) on page 205
- [Configuring VRRP on a V3 interface](#) on page 211

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 VSP Operating System Software*, NN47227-508.

For more information about Routing Information Protocol (RIP) and Open Shortest Path First (OSPF), see *Configuring OSPF and RIP on Avaya Virtual Services Platform 4000 Series*, NN46251-506.

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 4000.

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

Table continues...


Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
9	255.255.255.128	510	126
10	255.255.255.192	1 022	62
11	255.255.255.224	2 046	30
12	255.255.255.240	4 094	14
13	255.255.255.248	8 190	6
14	255.255.255.252	16 382	2
Class C			
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

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).



9577EA

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.

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 the following figure, 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.

*** Note:**

You can configure only one CLIP interface with an IPv6 address, which can be only used as a source IPV6 address for IS-IS.

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 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.

IPv6 Circuitless IP

IPv6 Circuitless IP (CLIP) is a virtual interface that is not associated with any physical port. You can use an IPv6 CLIP interface to provide uninterrupted connectivity to your switch as long as an actual path exists to reach the device. The system treats the IPv6 CLIP interface like an IPv6 interface and treats the network associated with the IPv6 CLIP as a local network attached to the device. This route always exists and the circuit is always up because no physical attachment exists.

You can use an IPv6 CLIP address as a logical IPv6 address for network management, as well as for other purposes. The IPv6 CLIP is typically a host address with any prefix length. You can redistribute this address as part of any other routing protocol update, so that the CLIP address is known to neighbors and available for use in routing or other types of connectivity. You can use IPv6 CLIP for many kinds of management connectivity such as telnet or SSH. You can also use IPv6 CLIP as a source IP address for sending Syslog messages.

IPv6 CLIP restrictions and limitations

This section describes the restrictions and limitations associated with IPv6 CLIP.

- Stateless address autoconfiguration (SLAAC) is not supported on IPv6 CLIP interfaces.
- IPv6 CLIP does not support link-local address configuration.
- To configure an IPv6 address with a prefix length from 65 to 127 on a CLIP interface, you must enable the `ipv6` mode flag.

*** Note:**

This limitation applies only to VSP 8000 switches. It does not apply to VSP 4000 switches.

- Neighbor discovery (ND) does not run on an IPv6 CLIP interface. Therefore, the system does not detect duplicate IPv6 address assignment to this interface.
- Multiple IPv6 address configuration on an IPv6 CLIP interface is not supported.
- You can configure a maximum of 64 IPv6 CLIP interfaces.
- IPv6 CLIP interface is enabled by default and it cannot be disabled.
- You cannot configure an IPv6 CLIP interface as the source or destination endpoint of an IPv6-in-IPv4 tunnel.

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 4000 Series 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 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. Router ports are VLANs that route IP packets and bridge nonroutable traffic in a single-port VLAN.

Virtual routing between VLANs

Virtual Services Platform 4000 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.

Router ports

Virtual Services Platform 4000 also supports router ports. A router port is a single-port VLAN that routes IP packets and bridges all nonroutable traffic. The difference between a router port and a standard IP protocol-based VLAN configured to route traffic is that the routing interface of the router port is not subject to the spanning tree state of the port. A router 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 router port is a single-port VLAN, each router port decreases the number of available VLANs by one and uses one VLAN ID.

Virtual Services Platform 4000 allows ip routing to be enabled on a maximum of 256 VLANs/router ports.

Equal Cost Multipath

With Equal Cost Multipath (ECMP), Avaya Virtual Services Platform 4000 Series can determine up to four 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:

*** Note:**

ECMP is supported on both Global Routing Table (GRT) and Virtual routing and forwarding (VRF).

- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Border Gateway Protocol (BGP)
- Virtual routing and forwarding (VRF)
- 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 4000 Series 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 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 4000 Series 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 4000 Series, 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 reenables the route.

! Important:

Changing route preferences is a process-oriented operation that can affect system performance and network reachability. Therefore, Avaya recommends that if you want to change preferences for static routes or routing protocols, you do so when you configure routes or before you enable routing protocols.

On Virtual Services Platform 4000, 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
Interior BGP	175
Staticv6	5
OSPFv3 intra-area	20
OSPFv3 inter-area	25

Table continues...

Protocol	Default preference
OSPFv3 external type 1	120
OSPFv3 external type 2	125
SPBM level 1	7
SPBMv6 level 1	7

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.

*** Note:**

VSP 4000 supports a maximum of 256 IPv6 ingress port/vlan security ACL/filters. IPv6 ingress QoS ACL/Filters and IPv6 Egress Security and QoS ACL/Filters are not supported.

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

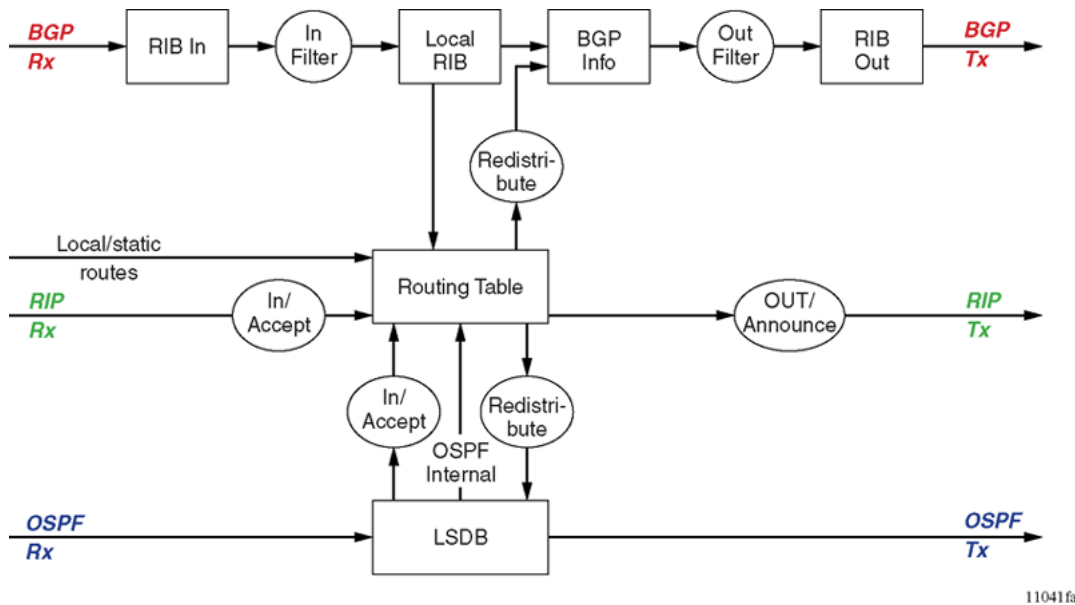


Figure 5: Route filtering for BGP, RIP, and OSPF routing protocols

The following figure shows how filters apply to the IS-IS protocol for Avaya 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.



10531eb

Figure 7: Route filtering stages

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

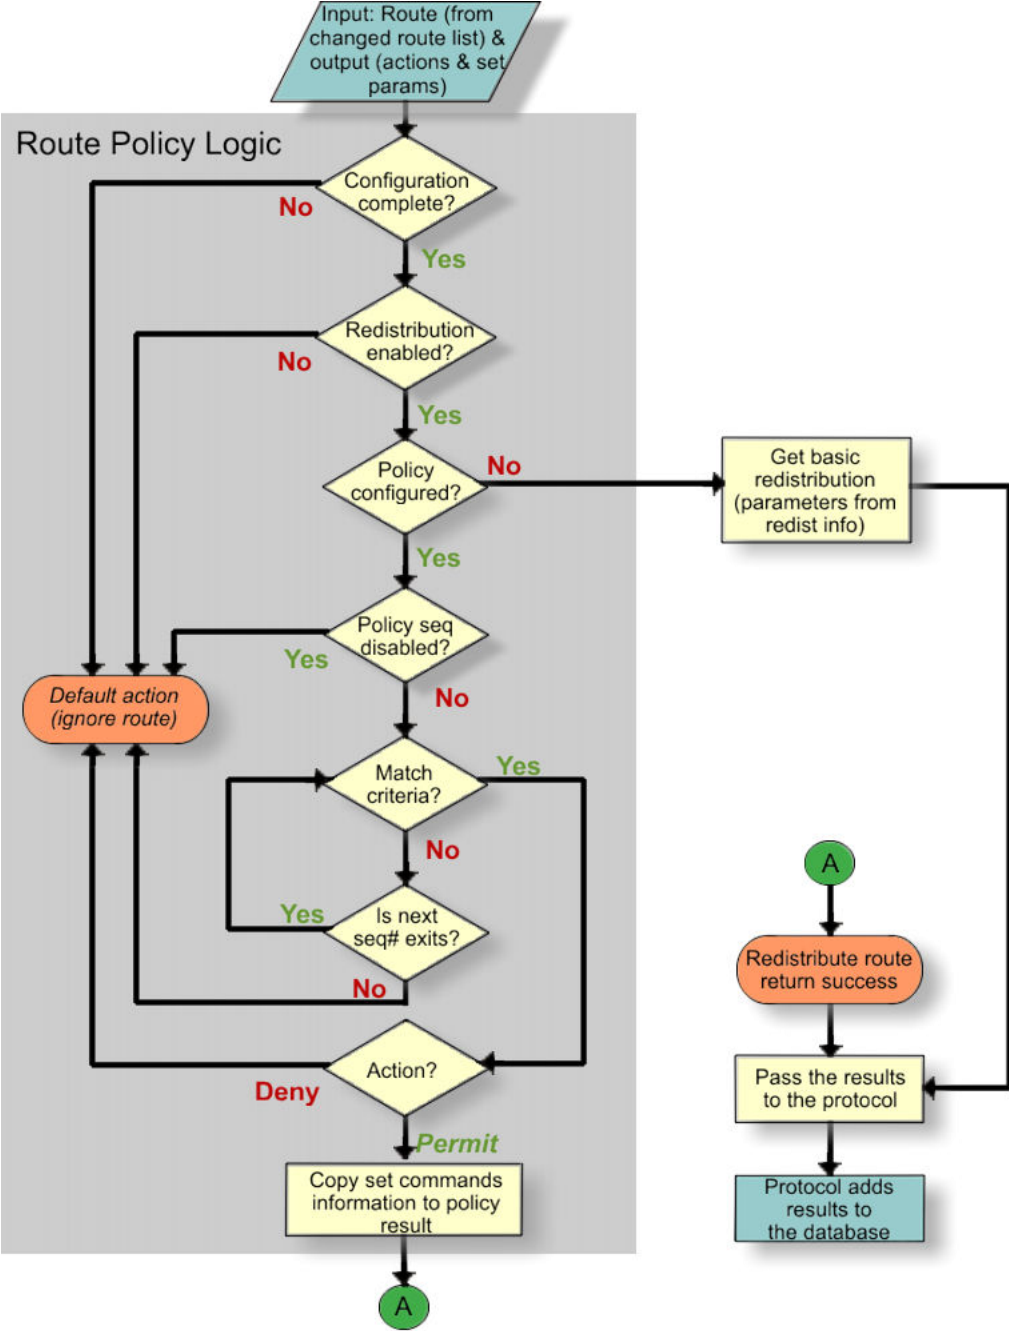


Figure 8: Route filtering logic

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.

* Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

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		

Table continues...

	Announce				Accept
	OSPF	Direct	RIP	BGP	RIP
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					
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

Table continues...

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
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
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

Table continues...

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
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					
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

Table continues...

	Redistribute			Accept	Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
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					
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-egg-as					
Set Tag					
Set Weight				Yes	

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 4000 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 4000 Series is between the two hosts. To enable communication between the two hosts, the Avaya Virtual Services Platform 4000 Series responds to the ARP request with the IP address of host C but with its own MAC address.

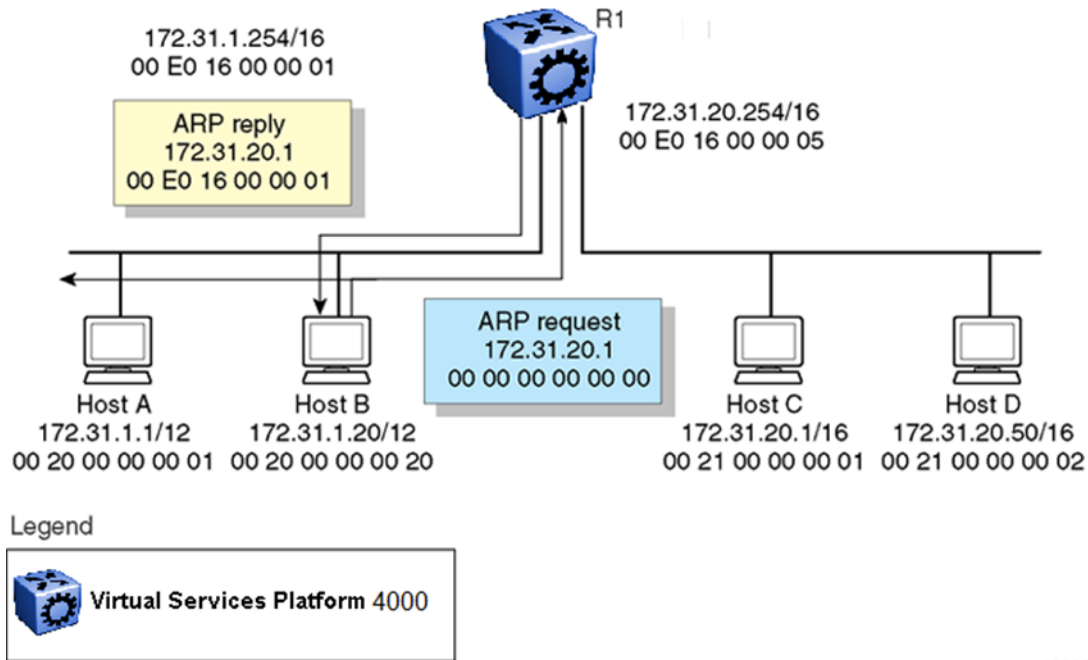


Figure 9: Proxy ARP operation

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 10: 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 11: 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 4000 Series, 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 4000 Series 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.

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.

Note:

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

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 12: Virtual Router Redundancy Protocol configuration

The switch supports 253 VRRP interfaces for each VRF and 253 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
- 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 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, 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, the switch 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.

* Note:

In this context, *local* implies an address from the same VRF as the IP interface where VRRP is being configured.

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 vIST

towards the master VRRP router. In this case, the vIST 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.

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 communicate with other switch ports and Avaya Networking products, such as ERS 8800 and VSP 9000, with the same settings.

VRRP guidelines

VRRP guidelines

VRRP provides another layer of resiliency to your network design by providing default gateway redundancy for end users. If a VRRP-enabled router that connects to the default gateway fails, failover to the VRRP backup router ensures no interruption for end users who attempt to route from their local subnet.

Only the VRRP Master router forwards traffic for a given subnet. The backup VRRP router does not route traffic destined for the default gateway.

To allow both VRRP switches to route traffic, VOSS has an extension to VRRP, the BackupMaster, that creates an active-active environment for routing. If you enable BackupMaster on the backup router, the backup router no longer switches traffic to the VRRP Master. Instead the BackupMaster routes all traffic received on the BackupMaster IP interface according to the switch routing table.

Figure 13: VRRP with BackupMaster

Avaya recommends that you stagger VRRP instances on a network or subnet basis. The following figure shows the VRRP Masters and BackupMasters for two subnets. For more information about how to configure VRRP using the Avaya Command Line Interface (ACLI) and Enterprise Device Manager (EDM), see [VRRP configuration using ACLI](#) on page 191 and [VRRP configuration using EDM](#) on page 207.

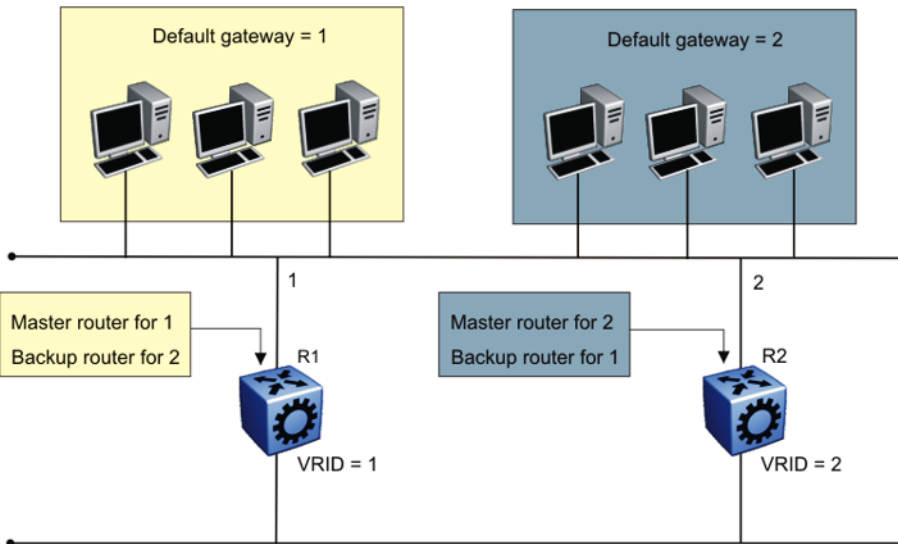


Figure 14: VRRP network configuration

The VRRP BackupMaster uses the VRRP standardized backup switch state machine. Thus, VRRP BackupMaster is compatible with standard VRRP.

Avaya recommends that you use the following best practices to implement VRRP:

- Do not configure the virtual address as a physical interface that is used on the routing switches. Instead, use a third address, for example:
 - Interface IP address of VLAN A on Switch 1 = x.x.x.2
 - Interface IP address of VLAN A on Switch 2 = x.x.x.3
 - Virtual IP address of VLAN A = x.x.x.1

*** Note:**

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

- Configure the VRRP holddown timer with enough time that the Interior Gateway Protocol (IGP) routing protocol has time to update the routing table. In some cases, configuring the VRRP holddown timer to a minimum of 1.5 times the IGP convergence time is sufficient. For OSPF, Avaya recommends that you use a value of 90 seconds if you use the default OSPF timers.
- Implement VRRP BackupMaster for an active-active configuration (BackupMaster works across multiple switches that participate in the same VRRP domain).

- Configure VRRP priority as 200 to configure VRRP Master.
- Stagger VRRP Masters between switches in the core to balance the load between switches.
- If you implement VRRP Fast, you create additional control traffic on the network and also create a greater load on the CPU. To reduce the convergence time of VRRP, the VRRP Fast feature allows the modification of VRRP timers to achieve subsecond failover of VRRP. Without VRRP Fast, normal convergence time is approximately 3 seconds.
- Do not use VRRP BackupMaster and critical IP at the same time. Use one or the other.

VRRP and spanning tree

VOSS can use one of two spanning tree protocols: Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP).

VRRP protects clients and servers from link or aggregation switch failures. Configure the network to limit the amount of time a link is out of service during VRRP convergence. The following figure shows two possible configurations of VRRP and spanning tree; configuration A is optimal and configuration B is not.




Figure 15: VRRP and STG configurations

In this figure, configuration A is optimal because VRRP convergence occurs within 2 to 3 seconds. In configuration A, three spanning tree instances exist and VRRP runs on the link between the two routers. Spanning tree instance 2 exists on the link between the two routers, which separates the link between the two routers from the spanning tree instances found on the other devices. All uplinks are active.

In configuration B, VRRP convergence takes between 30 and 45 seconds because it depends on spanning tree convergence. After initial convergence, spanning tree blocks one link (an uplink), so

only one uplink is used. If an error occurs on the uplink, spanning tree reconverges, which can take up to 45 seconds. After spanning tree reconvergence, VRRP can take a few more seconds to fail over.

VRRP and ICMP redirect messages

You can use VRRP and Internet Control Message Protocol (ICMP) together. However, doing so can provide nonoptimal network performance.

Consider the network shown in the following figure. Traffic from the client on subnet 30.30.30.0, destined for the 10.10.10.0 subnet, is sent to routing switch 1 (VRRP Master). Routing switch 1 forwards this traffic on the same subnet to routing switch 2, where it is routed to the destination. With ICMP redirect enabled, for each packet received, routing switch 1 sends an ICMP redirect message to the client to inform it of a shorter path to the destination through routing switch 2.




Figure 16: ICMP redirect messages

If network clients do not recognize ICMP redirect messages, disable ICMP redirect messages on the VOSS switch to avoid excessive ICMP redirect messages. Avaya recommends the network design shown in the following figure.

Ensure that the routing path to the destination through both routing switches has the same metric to the destination. One hop goes from 30.30.30.0 to 10.10.10.0 through routing switch 1 and routing switch 2.




Figure 17: Avoiding excessive ICMP redirect messages without SMLT

VRRPv3

VRRPv3 is a combined protocol for both IPv4 and IPv6. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IPv4 or IPv6 addresses associated with a virtual router is called Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Backups wait for a Master and take ownership when the Master is no longer detected.

The election protocol provides dynamic failover in the forwarding responsibility when the Master is unavailable. VRRP for IPv4 gains a higher-availability default path without configuring dynamic routing or router discovery protocols on every end-host. VRRP for IPv6 gains a quick switch-over to Backup routers compared to the standard IPv6 Neighbor Discovery mechanisms.

In the previous releases VRRPv2 for IPv4 and VRRPv3 for IPv6 (draft-ietf-vrrp-ipv6-spec-08.txt) were implemented on the switch. The current release adds VRRPv3 for IPv4 and makes VRRPv3 for IPv6 compliant to RFC5798. The existing VRRPv2 for IPv4 feature remains unchanged.

VRRPv3 guidelines

The switch also supports VRRPv2 for IPv4. If VRRP IPv6 is set on an interface it will run independent of the IPv4 version. Set the version of the VRRP IPv4 on the interface before configuring any other IPv4 VRRP attributes. The version is not set by default to a particular value. However, when sourcing older configuration files that do not have the version saved on them, the router will set it to VRRPv2 by default. If you change the version, all IPv4 configuration under that interface is automatically removed, and you are prompted for a confirmation before this operation.

ACL configuration is done through `ip vrrp` or `ipv6 vrrp` nodes; CLI commands for IPv4 will be common for version 2 and version 3 and CLI commands for IPv6 will remain the same.

Following features are added to make both IPv4 and IPv6 VRRPv3 features compliant to RFC5798:

- Advertisement vs Fast-advertisement — Prior to RFC5798, the minimum advertisement interval was 1 second, with Fast-advertisement a sub-second interval could be configured. When this feature is enabled, the VRRP ADVERTISEMENT packets are sent with type 7 instead of 1. With RFC5798 the sub-second interval is standardised, and in this release we send all packets for VRRPv3 with type 1. The use of Fast-advertisement remains the same. VRRPv2 packets will still be sent with type 7, if Fast-advertisement is enabled.
- Add Master-advertisement-interval — Previously, all virtual routers on the same VLAN were to have the same Advertisement-Interval configured. RFC 5798 states that, you can have different Advertisement Intervals on the Master and Backup. On Master, the Master-advertisement-interval and the Advertisement-Interval have the same value. On Backup, the Master-advertisement-interval is used to calculate the timers, and the locally configured Advertisement-Interval is ignored until the Backup transitions to Master. The Master-advertisement-interval value is put in the advertisement packet type sent by the Master
- Transition to master as specified in RFC 5798 — Previously, if a Backup receives an advertisement with a lower priority (or same priority but lower IP), it immediately sends its own advertisement and transitions to Master. However, RFC 5798 states that such packet must be discarded, which means it will transition to Master after the Master_Down_Timer expires
- Add skew-time — Skew-time is a new feature supported in this release, RFC 5798 states that skew-time is calculated depending on the priority, and Master-advertisement-interval assures that the Backup with highest priority sends the first advertisement when the Master goes down
Skew time is calculated using the formula: $((256 - \text{priority}) * \text{Master_Adver_Interval}) / 256$.
- Add preempt-mode — Preempt-mode is a new attribute, and it is different from the `ipv6 vrrp <vrid> action preempt` command, which is an operational command issued when you want to stop the hold-down timer. RFC5798 states that preempt-mode should be set to false when you do not want a higher priority Backup to transition to Master. By default, it is set to true

*** Note:**

Accept-mode is not fully implemented for IPv4 VRRPv3. You can only ping the virtual IP address, the same way as it is for IPv4 VRRPv2.

Layer 3 switch clustering and multicast SMLT

Switch clustering is the logical aggregation of two nodes to form one logical entity known as the switch cluster. The two peer nodes in a switch cluster connect using a virtual interswitch trunk (vIST). The vIST exchanges forwarding and routing information between the two peer nodes in the cluster. This section provides guidelines for switch clusters that use multicast and Split Multilink Trunking (SMLT).

General guidelines

The following list identifies general guidelines to follow if you use multicast and switch clustering:

- Enable Protocol Independent Multicast - Sparse Mode (PIM-SM) on the vIST VLAN for fast recovery of multicast. A unicast routing protocol is not required.
- Enable Internet Group Management Protocol (IGMP) snooping and proxy on the edge switches.

The following figure shows multicast behavior in an SMLT environment. The configuration in the following figure provides fast failover if the switch or rendezvous point (RP) fails.





Figure 18: Multicast behavior in SMLT environment

In Multicast behavior in SMLT environment the following actions occur:

1. The multicast server sends multicast data towards the source designated router (DR).
2. The source DR sends register messages with encapsulated multicast data towards the RP.
3. After the client sends IGMP membership reports towards the multicast router, the router creates a (*,G) entry.

4. The RP sends join messages towards the source DR on the reverse path.
5. After the source DR receives the join messages, it sends native multicast traffic.
6. After SW_B or SW_D receives multicast traffic from upstream, it forwards the traffic on the vIST as well as on the SMLT link. Other aggregation switches drop multicast traffic received over the vIST at egress. This action provides fast failover for multicast traffic. Both SW_D and SW_E (Aggregation switches) have similar (S,G) records.
7. In case of SW_D or RP failure, SW_B changes only the next-hop interface towards SW_E. Because the circuitless IP (CLIP) RP address is the same, SW_B does not flush (S,G) entries and achieves fast failover.

Multicast triangle topology

A triangle design is an SMLT configuration that connects edge switches or SMLT clients to two aggregation switches. Connect the aggregation switches together with a vIST that carries all the SMLT trunks configured on the switches.

VOSS switches support the following triangle configurations:

- a configuration with Layer 3 PIM-SM routing on both the edge and aggregation switches
- a configuration with Layer 2 snooping on the client switches and Layer 3 routing with PIM-SM on the aggregation switches

To avoid using an external query device to provide correct handling and routing of multicast traffic to the rest of the network, use the triangle design with IGMP Snoop at the client switches. Use multicast routing at the aggregation switches as shown in the following figure.




Figure 19: Multicast routing using PIM-SM

Client switches run IGMP Snoop or PIM-SM, and the aggregation switches run PIM-SM. This design is simple and, for the rest of the network, PIM-SM performs IP multicast routing. The aggregation

switches are the query devices for IGMP, so an external query device is not required to activate IGMP membership. These switches also act as redundant switches for IP multicast.

Multicast data flows through the vIST link when receivers are learned on the client switch and senders are located on the aggregation switches, or when sourced data comes through the aggregation switches. This data is destined for potential receivers attached to the other side of the vIST. The data does not reach the client switches through the two aggregation switches because only the originating switch forwards the data to the client switch receivers.

*** Note:**

Always place multicast receivers and senders on the core switches on VLANs different from those that span the vIST.

The following figure shows a switch clustering configuration with a single switch cluster core and dual-connected edge devices. This topology represents different VLANs spanning from each edge device and those VLANs routed at the switch cluster core. You can configure multiple VLANs on the edge devices, 802.1Q tagged to the switch cluster core.




Figure 20: Multicast SMLT triangle

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either Virtual Router Redundancy Protocol (VRRP) BackupMaster or Routed SMLT (RSMLT) Layer 2 Edge on the switch cluster core.

Square and full-mesh topology multicast guidelines

A square design connects a pair of aggregation switches to another pair of aggregation switches. A square design becomes a full-mesh design if the aggregation switches are connected in a full-mesh. The VOSS switch supports Layer 3 IP multicast (PIM-SM only) over a full-mesh SMLT or RSMLT configuration.

In a square design, configure all switches with PIM-SM. Place the bootstrap router (BSR) and RP in one of the four core switches; Avaya recommends that you place the RP closest to the source. If using PIM-SM over a square or full-mesh configuration, enable the `multicast smlt-square` flag.

The following three figures show switch clustering configurations with two-switch cluster cores and dual-connected edge devices.

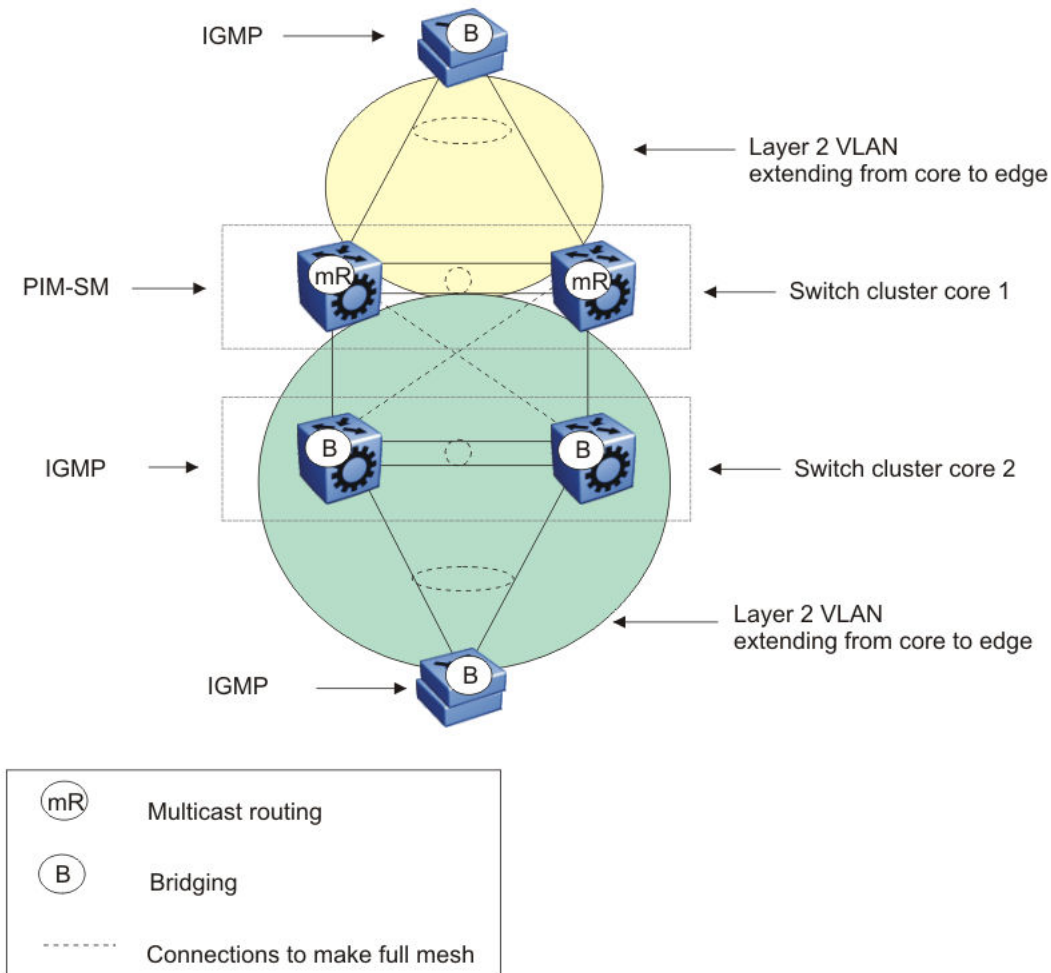


Figure 21: Multicast SMLT square 1

In the preceding figure, only one of the switch cluster cores performs Layer 3 multicast routing while the other is strictly Layer 2. Configure multiple VLANs on the edge devices, 802.1Q tagged to the switch cluster cores.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster core.

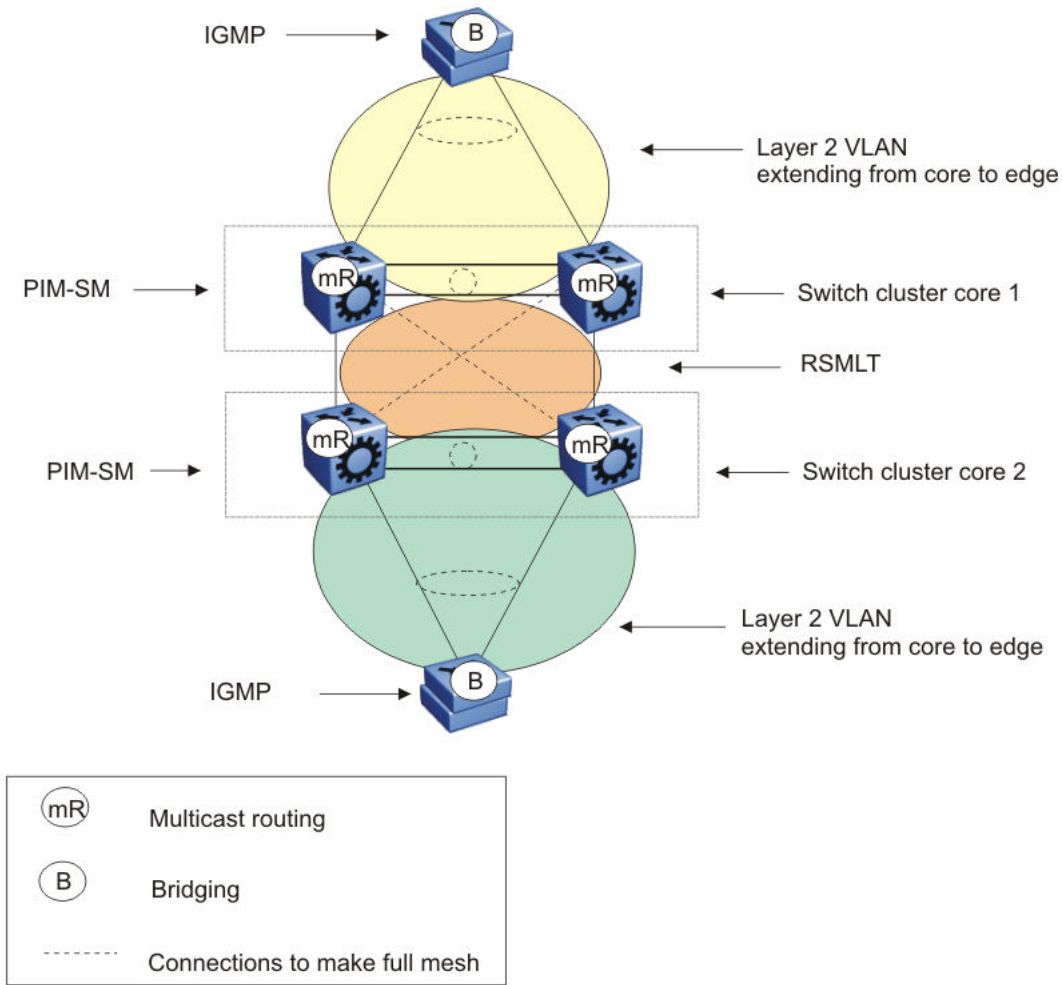


Figure 22: Multicast SMLT square 2

In the preceding figure, both of the switch cluster cores performs Layer 3 multicast routing, while the edge devices are Layer 2 IGMP.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster cores. Do not enable VRRP on the RSMLT VLAN between switch cluster cores.

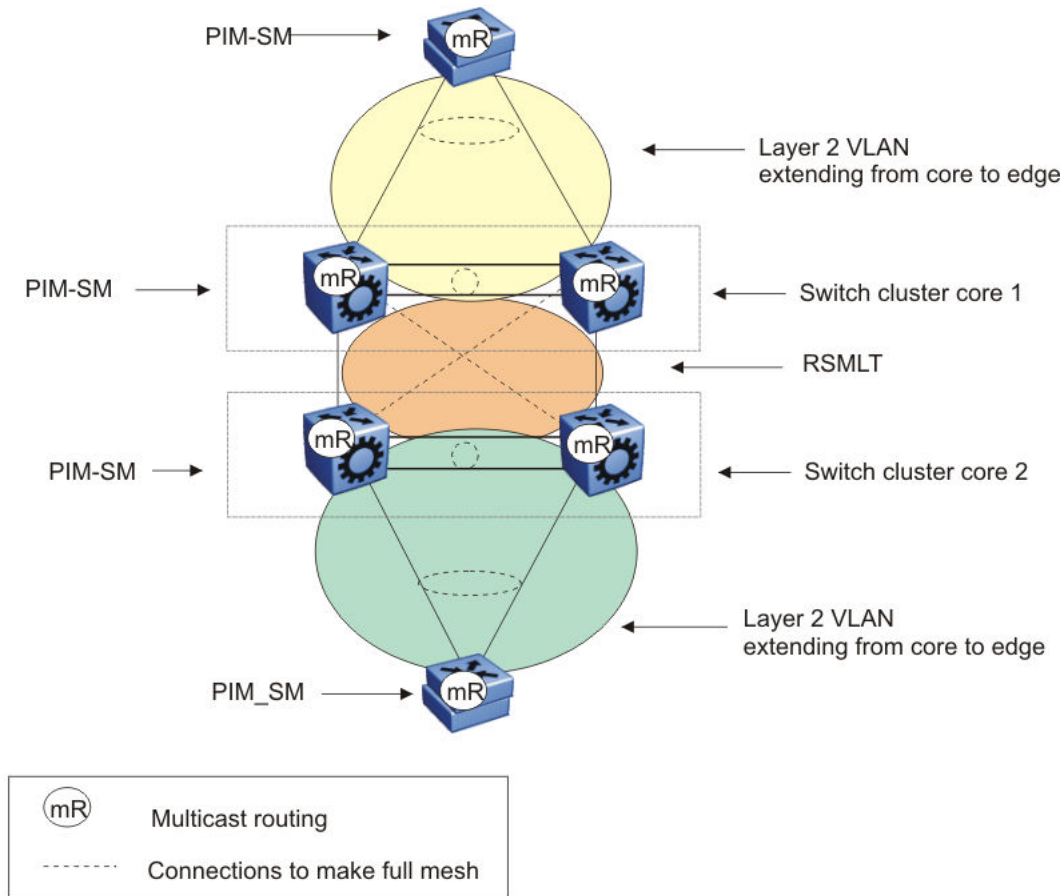


Figure 23: Multicast SMLT square 3

In the preceding figure, both of the switch cluster cores and the edge devices perform Layer 3 multicast routing.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster cores. Do not enable VRRP on the RSMLT VLAN between switch cluster cores.

SMLT and multicast traffic issues

If PIM-SM or other multicast protocols are used in an SMLT environment, enable the protocol on the vIST. Routing protocols in general are not run over an vIST but multicast routing protocols are an exception. When using PIM-SM and a unicast routing protocol, ensure the unicast route to the BSR and RP has PIM-SM active and enabled. If multiple OSPF paths exist and PIM-SM is not active on each pair, the BSR is learned on a path that does not have PIM-SM active. The following figure demonstrates this issue.




Figure 24: Unicast route example

The network configuration in the preceding figure is as follows:

- 5510A is on VLAN 101.
- 5510B is on VLAN 102.
- VSP Switch B is the BSR.
- VSP Switch A and VSP Switch B have OSPF enabled.
- PIM is enabled and active on VLAN 101.
- PIM is either disabled or passive on VLAN 102.

In this example, the unicast route table on VSP Switch A learns the BSR on VSP Switch B through VLAN 102 using OSPF. The BSR is either not learned or does not provide the RP to VSP Switch A.

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:

- 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 25: SMLT and RSMLT in Layer 3 environments](#) on page 58 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 vIST messaging) of their physical IP and 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 routes IP traffic that is directed to the physical MAC of the IP or the physical MAC of the peer IP on that VLAN, and when one of them is down the other 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

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 25: 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 to v1ST's L2VSN another routed VLAN with lower protocol metrics as a 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 v1ST's L2VSN. For more information about how to configure SMLT in a Layer 2 environment, see *Configuring Link Aggregation, MLT and SMLT on VSP Operating System Software*, NN47227-503.

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 v1ST pair is active, traffic forwarding is available for both next hops R1/R2 and R3/R4.

RSMLT edge support

The switch 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 vIST 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 vIST switches come back, the vIST 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 vIST 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).

Enable or disable IPv4 ICMP broadcast

On IPv4 networks, a packet can be directed to an individual machine or broadcast to an entire network. When a packet is sent to an IP broadcast address from a machine on the local network, that packet is delivered to all machines on that network.

If a packet that is broadcast is an ICMP echo request packet, the machines on the network receive this ICMP echo request packet and send an ICMP echo reply packet back. When all the machines on a network respond to this ICMP echo request, the result can be severe network congestion or outages.

VOSS devices always respond to the IPv4 ICMP packets sent to broadcast address. VOSS 5.1 introduces a feature to disable the processing of IPv4 ICMP packets sent to broadcast address. On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when the packets reach the control plane.

You can disable or enable the IPv4 ICMP broadcast processing at the VRF level.

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.

*** Note:**

The prompt for the non-PowerPlus chassis is VSP-4850GTS. The prompt for the PowerPlus chassis is VSP-4850GTS-PWR+. The prompt for the Fiber box is VSP-4450 GSX. For consistency, this document uses the VSP-4850GTS prompt.

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-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

```
VSP-4850GTS-PWR+:1(config)#interface vlan 200
```



```
VSP-4850GTS-PWR+: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-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1#configure terminal
VSP-4850GTS-PWR+:1(config)#interface vlan 200
VSP-4850GTS-PWR+:1(config-if)#ip arp-proxy enable
```

Showing ARP information

Before you begin

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

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 4000 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-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1#configure terminal
VSP-4850GTS-PWR+:1#interface vlan 200
VSP-4850GTS-PWR+:1(config-if)#show ip arp interface
```

```
=====
                                     Port Arp
=====
PORT_NUM  DOPROXY   DORESP
-----
1/1       false     true
1/2       false     true
1/3       false     true
1/4       false     true
1/5       false     true
1/6       false     true
1/7       false     true
1/8       false     true
1/9       false     true
1/10      false     true
1/11      false     true
1/12      false     true
1/13      false     true
1/14      false     true
1/15      false     true
1/16      false     true
1/17      false     true

--More-- (q = quit)
```

Variable definitions

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

Table 7: 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 router port.
interface	Displays ARP interface configuration information.
-s	Specifies a subnet. You must indicate the IP address followed by the subnet mask expressed as <A.B.C.D> <A.B.C.D>.
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<0-16>	Specifies a VRF name expressed as text from 0 to 16 characters in length.

Table continues...

Variable	Value
	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 <i>WORD</i> <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 8: 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 9: 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

Before you begin

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

About this task

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.

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

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Add ARP entries:

```
VSP-4850GTS-PWR+:1(config)#ip arp 46.17.159.128 00-16-76-7D-80-C2 1/1
```

Configure static multicast MAC entries:

```
VSP-4850GTS-PWR+:1(config)#ip arp static-mcast 46.17.159.128
00-16-76-7D-80-C2 vid 200
```


Variable definitions

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

Table 10: Variable definitions

Variable	Value
multicast-mac-flooding [enable]	<p>Determines whether ARP entries for multicast MAC addresses are associated with the VLAN or the port interface on which they were learned.</p> <p>Use the <code>no</code> operator to delete a static entry from the ARP table: <code>no ip arp multicast-mac-flooding [enable]</code></p> <p>To configure this option to the default value, use the default operator with this command.</p>
request-threshold <50-1000>	<p>Configures the maximum number of outstanding ARP requests that a device can generate. The range is 50–1000. The default value is 500.</p> <p>To configure this option to the default value, use the default operator with this command.</p>
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.

Table continues...

Variable	Value
	<p>To configure this option to the default value, use the default operator with this command.</p> <p> Note:</p> <p>The aging of ARP records is tied to the aging of MAC records. The ARP record for a given IP address is not removed unless the associated MAC record ages out and the router stops receiving a response to ARP requests for that IP address. In cases where the ARP aging time is set to less than the MAC aging time, the switch waits until the MAC ages out before deleting the ARP for an inactive host.</p>
<A.B.C.D>	Adds ARP entries.

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-4850GTS-PWR+:1>enable
```

Clear ARP entries:

```
VSP-4850GTS-PWR+:1#clear ip arp interface gigabitethernet 1/16
```

Variable definitions

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

Table 11: 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

Before you begin

- You must log on to one of the following ACLI modes:
 - Privileged EXEC
 - VRF Router Configuration

About this task

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

When you use the interface parameter with the `show ip arp` command you can display ARP configuration information only for a specific Virtual Services Platform 4000 switch.

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

Procedure

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] [-s | vrf | vrfids
| <A.B.C.D>] [vlan <1-4084>] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-4850GTS:1#show ip arp
```

```
=====
                        IP Arp - GlobalRouter
=====
IP_ADDRESS      MAC_ADDRESS      VLAN    PORT    TYPE      TTL(10 Sec)  TUNNEL
-----
47.17.41.1      00:11:f9:5b:10:42 4000    1/46    DYNAMIC   1929
47.17.41.11     cc:f9:54:ae:9c:80 4000    1/46    DYNAMIC   1938
47.17.41.114    00:1c:c4:d6:28:ba 4000    1/46    DYNAMIC   2043
47.17.41.255    ff:ff:ff:ff:ff:ff 4000    -        LOCAL     2160
=====
                        IP Arp Extn - GlobalRouter
=====
```

```
MULTICAST-MAC-FLOODING    AGING (Minutes)    ARP-THRESHOLD
-----
disable                    360                500

25 out of 25 ARP entries displayed

VSP-4850GTS:1#
```

Variable definitions

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

Table 12: 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 router port.
interface	<p>Displays ARP interface configuration information.</p> <p>Use the following parameters to display ARP table information specifically for:</p> <ul style="list-style-type: none"> gigabitEthernet {slot/port[-slot/port][,...]} displays IP ARP gigabitEthernet interface information VLAN <1-4084> displays IP ARP VLAN interface information <p>Example: <code>show ip arp interface vlan 1</code></p>
spbm-tunnel-as-mac	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
vlan	<p>Displays ARP entries for a particular VLAN ID in a range from 1 to 4084.</p> <p>Use these parameters to display ARP table information specifically for:</p> <ul style="list-style-type: none"> vrf WORD<0-16>—the VLAN VRF name in a range from 0 to 16 characters vrfids WORD<0-512>—the VLAN VRF ID in a range from 0 to 512 <p>Example: <code>show ip arp vlan 1 vrf 1</code></p>
vrf WORD <0-16>	<p>Specifies the name of the VRF.</p> <p>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.</p>
vrfids WORD <0-512>	Specifies the VRF ID.

Table continues...

Variable	Value
	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.

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

Table 13: 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.
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.

Configuring Gratuitous ARP

Use the following procedure to configure Gratuitous Address Resolution Protocol (ARP). When Gratuitous ARP is enabled the switch allows all Gratuitous ARP request packets. The default is enabled.

If you disable Gratuitous ARP, the switch only allows Gratuitous ARP packets associated with Routed Split Multi-Link Trunking (RSMLT) or Virtual Router Redundancy Protocol (VRRP), and the switch discards all other Gratuitous ARP request packets.

About this task

ARP translates network layer (layer 3) IP addresses into link layer (layer 2) MAC addresses. A host sends a Gratuitous ARP request packet to inform other hosts of the existence of an interface on the network, so other local hosts can update their ARP tables. If the IP or MAC address changes, or in the event of a failover, a host sends a Gratuitous ARP request packet to inform other hosts to update their ARP tables.

VRRP and RSMLT use gratuitous ARP to update the MAC address tables on switches.

Procedure

1. Enter Global Configuration mode:

```
enable  
configure terminal
```

2. Enable Gratuitous ARP:

```
ip gratuitous-arp
```

3. **(Optional)** Disable Gratuitous ARP:

```
no ip gratuitous-arp
```

4. **(Optional)** Configure Gratuitous ARP to the default value:

```
default ip gratuitous-arp
```

5. Save the changed configuration.

```
save config [backup WORD<1-99>][file WORD<1-99>][verbose]
```

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 router 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 > Basic**.
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. 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.

Table continues...

Name	Description
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.

Viewing and managing ARP

About this task

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.

ARP field descriptions

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

Name	Description
NetAddress	Specifies the IP address corresponding to the media-dependent physical address.
IfIndex	Identifies the router interface for this ARP entry: <ul style="list-style-type: none"> • Brouter interfaces are identified by the slot/port number of the brouter port. • VLAN interfaces are identified by the vlan name.
PhysAddress	Specifies the media-dependent physical address (that is, the Ethernet address).
Type	Specifies the type of ARP entry: <ul style="list-style-type: none"> • local—a locally configured ARP entry • static—a statically configured ARP entry • dynamic—a learned ARP entry
TimeToLive	Indicates the time to live where the ARP is configured.
DestIfIndex	Indicates the slot/port on which the ARP entry was learned. For brouter interfaces this is the same value as IfIndex, but for VLAN interfaces, it designates the particular port in the VLAN on which the ARP was learned.
DestVlanId	VLAN ID where the ARP is configured.
BMac	Identifies the backbone MAC address if the entry is learned from an SPBM network.
DestCvid	Identifies the customer VLAN ID for a Switched UNI port.

Creating static ARP entries

About this task

Use the following procedure to create a static ARP entry.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **IP**.
3. Click the **ARP** tab.
4. Click **Insert**.
5. In the **NetAddress** field, type the IP address.
6. Click **Port**.
OR
Click **Port in VLAN**
7. In the dialog box, select the interface.
8. Click **OK**.
9. In the **PhysAddress** field, type the MAC address.
10. Click **Insert**.

Configuring ARP proxy

About this task

With an ARP proxy, the Avaya Virtual Services Platform 4000 Series 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 > Basic**.
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 4000.

*** Note:**

The prompt for the non-PowerPlus chassis is VSP-4850GTS. The prompt for the PowerPlus chassis is VSP-4850GTS-PWR+. The prompt for the Fiber box is VSP-4450 GSX. For consistency, this document uses the VSP-4850GTS prompt.

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-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create the forwarding path from the client to the server:

```
VSP-4850GTS-PWR+: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-4850GTS-PWR+: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-4850GTS-PWR+: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 14: Variable definitions

Variable	Value
<code>fwd-path <A.B.C.D> <A.B.C.D></code>	<p>Configures the forwarding path from the client to the server.</p> <p>A.B.C.D is the IP address configured on an interface (a locally configured IP address) to forward or relay BootP or DHCP.</p> <p>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.</p> <p>Use the <code>no</code> operator to delete the forwarding path from the client to the server: <code>no ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D></code>.</p>
<code>fwd-path <A.B.C.D> <A.B.C.D> disable</code>	<p>Disables DHCP relaying on the path from the IP address to the server. This feature is disabled by default.</p> <p>A.B.C.D is the IP address configured on an interface (a locally configured IP address).</p> <p>A.B.C.D is the IP address of the DHCP server in the network.</p>
<code>fwd-path <A.B.C.D> <A.B.C.D> enable</code>	<p>Enables DHCP relaying on the path from the IP address to the server.</p> <p>A.B.C.D is the IP address configured on an interface (a locally configured IP address).</p>

Table continues...

Variable	Value
	<p>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.</p> <p>Use the no operator to disable DHCP: <code>no dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> enable</code>.</p> <p>To configure this option to the default value, use the default operator with the <code>ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D></code> command.</p>
<code>fwd-path <A.B.C.D> <A.B.C.D> mode <bootp bootp_dhcp dhcp></code>	<p>Modifies DHCP mode to forward BootP messages only, DHCP messages only, or both. The default is both.</p> <p>mode is {bootp bootp_dhcp dhcp}.</p>

Showing DHCP relay information

Before you begin

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

About this task

Display relay information to show relay information about DHCP routes and counters. For DHCP relay only 128 entries are supported.

Procedure

- Display information about DHCP relay forward paths:

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

- Display information about DHCP relay counters:

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

- Display the options for each listed interface:

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

Example

```
VSP-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSP-4850GTS-PWR+:1#show ip dhcp-relay interface
```

```
=====
Port Dhcp
=====
PORT VRF          MAX MIN          ALWAYS CIRCUIT REMOTE TRUST
NUM  NAME          ENABLE HOP SEC        BCAST  ID      ID      CIRC
=====
```

```

=====
                                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 15: Variable definitions

Variable	Value
vrf WORD<0-16>	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<0-16>]	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

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.

About this task

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.

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 <0-16>] [vrfids WORD <0-512>]
```

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

```
VSP-4850GTS-PWR+:1(config)#interface gigabitethernet 1/10
```

Enable the circuit ID:

```
VSP-4850GTS-PWR+:1(config-if)#ip dhcp-relay circuitID
```

Enable the remote ID:

```
VSP-4850GTS-PWR+:1(config-if)#ip dhcp-relay remoteID
```

Configure the circuit as trusted:

```
VSP-4850GTS-PWR+:1(config-if)#ip dhcp-relay trusted
```

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

```
VSP-4850GTS-PWR+: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 16: 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 <0–16>] [vrfids WORD <0–512>]` command.

Variable	Value
vrf WORD <0–16>	Displays DHCP counters for a particular VRF. WORD <0–16> 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

Before you begin

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

About this task

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

Procedure

Enable DHCP parameters on a specified port or VLAN:

```
ip dhcp-relay
```

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

```
VSP-4850GTS-PWR+:1(config)#interface gigabitethernet 1/10
```

Enable DHCP parameters on a specified port or VLAN:

```
VSP-4850GTS-PWR+: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 17: Variable definitions

Variable	Value
broadcast	<p>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: <code>no ip dhcp-relay broadcast</code>.</p> <p>To configure this option to the default value, use the default operator with this command.</p>
fwd-path <A.B.C.D> [vrid <1-255>]	<p>Creates a forward path server with a virtual router ID (or VRRP ID), a mode, and a state.</p> <p>A.B.C.D is the IP address.</p> <p>vrid <1-255> is the ID of the virtual router and is an integer from 1 to 255.</p> <p>Use the no operator to delete a forward path server with a specific value and virtual router ID: <code>no ip dhcp-relay fwd-path <A.B.C.D> [vrid <1-255>]</code></p> <p>To configure this option to the default value, use the default operator with this command.</p>
fwd-path <A.B.C.D> disable [vrid <1-255>]	<p>Disables a forward path server with a specific value and virtual router ID.</p> <p>A.B.C.D is the IP address.</p> <p>vrid <1-255> is the ID of the virtual router (or VRRP ID) and is an integer from 1 to 255.</p>
fwd-path <A.B.C.D> enable [vrid <1-255>]	<p>Enables a forward path server with a specific value and virtual router ID (or VRRP ID).</p> <p>A.B.C.D is the IP address in the form a.b.c.d.</p> <p>vrid <1-255> is the ID of the virtual router and is an integer from 1 to 255.</p>
fwd-path <A.B.C.D> mode <bootp bootp_dhcp dhcp> [vrid <1-255>]	<p>Configures the forward path mode for a VLAN. This command string is available only in VLAN Interface Configuration mode.</p> <p>A.B.C.D is the IP address in the form a.b.c.d.</p> <p>mode is a choice of bootp, dhcp, or bootp_dhcp.</p> <p>vrid <1-255> is the ID of the virtual router (or VRRP ID) and is an integer from 1 to 255.</p> <p>To configure this option to the default value, use the default operator with this command.</p>
max-hop <1-16>	<p>Configures the maximum number of hops before a BootP/DHCP packet is discarded (1 to 16). The default is 4.</p>

Table continues...

Variable	Value
	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. 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.

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

Before you begin

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

About this task

Configure UDP protocols to determine which UDP broadcasts are forwarded.

Procedure

1. Configure a UDP protocol:

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

2. Confirm your configuration:

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

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

```
VSP-4850GTS-PWR+: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 18: 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 <code>no ip forward-protocol udp <1-65535></code> .
[vrf WORD<0-16>]	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

Before you begin

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

About this task

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

Procedure

- Configure a UDP port forward entry:

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

- Confirm your configuration:

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

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Configure a UDP port forward entry:

```
VSP-4850GTS-PWR+: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 19: Variable definitions

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

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Configure the UDP port forwarding list:

```
VSP-4850GTS-PWR+:1(config)#ip forward-protocol udp portfwddlist 1
```

Log on to the VLAN interface configuration mode:

```
VSP-4850GTS-PWR+:1(config)#interface vlan 3
```

Configure the broadcast mask:

```
VSP-4850GTS-PWR+:1(config-if)#ip forward-protocol udp broadcastmask  
100.31.255.255
```

Configure the maximum time to live:

```
VSP-4850GTS-PWR+:1(config-if)#ip forward-protocol udp maxttl 10
```

Confirm the configuration:

```
VSP-4850GTS-PWR+:1(config-if)#show ip forward-protocol udp portfwddlist
```

Variable definitions

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

Table 20: Variable definitions

Variable	Value
<1-1000>	Creates a UDP port forwarding list in the range of 1–1000.
<1-65535> {A.B.C.D}	<p>Adds a UDP protocol port to the specified port forwarding list.</p> <p>1-65535 is a UDP protocol port in the range of 1–65535.</p> <p>A.B.C.D is an IP address in a.b.c.d format.</p> <p>Use the no operator to remove or delete a port forwarding list ID,</p> <pre>no ip forward-protocol udp portfwddlist <1-1000> <1-65535> <A.B.C.D>.</pre> <p>To configure this option to use the default value, use the default operator with this command.</p>
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}	<p>Configures the interface broadcast mask (the interface broadcast mask can be different from the interface mask).</p> <p>A.B.C.D is an IP address in a.b.c.d format.</p> <p>Use the no operator to delete the broadcast mask:</p> <pre>no ip forward-protocol udp broadcastmask {A.B.C.D}</pre> <p>To configure this option to the default value, use the default operator with this command.</p>
maxttl <1-16>	Configures the maximum time-to-live value (TTL) for the UDP broadcast forwarded by the interface. The range is 1–16.
portfwddlist <1-1000>	Assigns the list to the VLAN.
vlan <1-4084> [portfwddlist <1-1000>]	<p>Specifies the VLAN ID.</p> <p>If you use the portfwddlist variable with the vlan variable, it assigns the list to the specified VLAN, regardless of which VLAN context you currently configure.</p>

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. UDP forwarding only supports 128 entries.

There are four show options:

- 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<0-16>]
[vrfids WORD<0-512>]
```

2. Display the UDP port forwarding table:

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

Example

```
VSP-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+: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
```

Variable definitions

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

Table 21: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the IP address for the interface in a.b.c.d format.
vrf WORD<0–16>	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.

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 data from the following table in the DHCP Relay 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

- In the navigation tree, expand the following folders: **Configuration > VLAN**.
- Click **VLANs > Basic**.
- Select a VLAN.
- Click **IP**.
- Click the **DHCP Relay** tab.
- Select **Enable**.
- Configure the parameters as required.
- 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: <ul style="list-style-type: none"> bootp dhcp both

Table continues...

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 router 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. <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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.
Remoteld	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

About this task

The Avaya Virtual Services Platform 4000 Series configures the following protocols, by default:

- Time Service
- Terminal Access Controller Access Control System Plus (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 tree, expand the following folders: **Configuration > IP**.
2. Click **UDP Forwarding**.
3. Click **Insert**.
4. 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.

5. In the **Name** field, type a name for the protocol.
6. 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**.

- In the Insert Forwardings dialog box, select a destination UDP port from the defined protocols in the **DestPort** box.
- 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.

- 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: <ul style="list-style-type: none"> 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

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

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 to 100). 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>] [vrfs WORD<0-512>] [WORD <1-64>]
```

Example

Configure a prefix-list. Display the prefix list.

```
Switch> enable
Switch# configure terminal
Switch(config)# ip prefix-list LIST1 47.17.121.50/255.255.255.0
Switch(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: <ul style="list-style-type: none"> 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.

Table continues...

Variable	Value
	Lower bound and higher bound mask lengths together can define a range of networks.
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 an IPv6 prefix list

Use IPv6 prefix lists to allow or deny specific IPv6 route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create an IPv6 prefix list:

```
ipv6 prefix-list <WORD 1-64> <WORD 1-256> [<ge|le> <0-128>]
```

Use the same command to add additional prefixes to the list.

3. To rename the list:

```
ipv6 prefix-list <WORD 1-64> name <WORD 1-64>
```

Example

Create an IPv6 prefix list:

```
Switch:1<config># ipv6 prefix-list list4 4717:0:0:0:0:0:7933:6/64 ge 32
le 64
```

To rename the list:

```
Switch:1<config># ipv6 prefix-list list4 name list5
```

Variable definitions

Use the data in the following table to use the ipv6 prefix-list command..

Variable	Value
<WORD 1–256> [<ge le> <0–128>]	Creates or adds a prefix to the list. The default value is none. <ul style="list-style-type: none"> • <WORD 1–256> specifies the IP prefix and length. • <ge le> specifies greater than or equal to or less than or equal to. • <0–128> specifies the mask length in the range 0 to 128. <p>To disable this option, use no operator with the command</p>
name <WORD 1–64>	Names the prefix list. The default value is none.

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 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 Fabric Connect on VSP Operating System Software*, NN47227-510.

*** 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.

*** Note:**

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

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 22: Variable definitions

Variable	Value
as-path WORD<0-256>	<p>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.</p> <p><i>WORD</i> <0-256> specifies the list IDs of up to four AS-lists, separated by a comma.</p> <p>Use the no operator to disable match as-path: <code>no match as-path WORD<0-256></code></p>
community WORD<0-256>	<p>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.</p> <p><i>WORD</i> <0-256> specifies the list IDs of up to four defined community lists, separated by a comma.</p> <p>Use the no operator to disable match community: <code>no match community WORD<0-256></code></p>
community-exact enable	<p>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.</p> <p>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.</p> <p>enable enables match community-exact.</p> <p>Use the no operator to disable match community-exact: <code>no match community-exact enable</code></p>
interface WORD <0-259>	<p>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.</p>

Table continues...

Variable	Value
	<p><i>WORD</i> <0-259> specifies the name of up to four defined prefix lists, separated by a comma.</p> <p>Use the no operator to disable match-interface: <code>no match interface WORD <0-259></code></p>
metric <0-65535>	<p>Configures the device to match the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored.</p> <p><0-65535> specifies the metric value. The default is 0.</p>
network WORD <0-259>	<p>Configures the device to match the destination network against the contents of the specified prefix lists.</p> <p><i>WORD</i> <0-259> specifies the name of up to four defined prefix lists, separated by a comma.</p> <p>Use the no operator to disable match network: <code>no match network WORD <0-259></code>.</p>
next-hop WORD<0-259>	<p>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.</p> <p><i>WORD</i> <0-259> specifies the name of up to four defined prefix lists, separated by a comma.</p> <p>Use the no operator to disable match next hop: <code>no match next-hop WORD<0-259></code></p>
protocol WORD<0-60>	<p>Configures the device to match the protocol through which the route is learned.</p> <p><i>WORD</i> <0-60> is xxx, where xxx is local, ospf, ebgp, isis, rip, static, or a combination separated by , in a string length 0–60.</p> <p>Use the no operator to disable match protocol: <code>no match protocol WORD<0-60></code></p>
route-source WORD<0-259>	<p>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.</p> <p><i>WORD</i> <0-259> specifies the name of up to four defined prefix lists, separated by a comma.</p> <p>Use the no operator to disable match route source: <code>no match route-source WORD<0-259></code></p>
route-type {any local internal external external-1 external-2}	<p>Configures a specific route type to match (applies only to OSPF routes).</p>

Table continues...

Variable	Value
	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 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 23: 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: <code>no set as-path WORD<0-256></code>
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: <code>no set automatic-tag enable</code>
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.

Table continues...

Variable	Value
	Use the no operator to delete the community number: <code>no set community WORD<0-256></code>
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<0-1027>	Configures the device to replace the destination network of the route that matches this policy with the contents of the specified prefix list. WORD<0-1027> specifies one prefix list by name. Use the no operator to disable set injectlist: <code>no set injectlist</code>
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. <0-65535> 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: <code>no set mask</code>
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: <code>no set next-hop</code>
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: <code>no set nssa-pbit enable</code>

Table continues...

Variable	Value
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 24: 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 25: 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 Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya Fabric Connect on VSP Operating System Software*, NN47227-510.

* Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

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 26: Variable definitions

Variable	Value
<A.B.C.D>	Specifies the IP address.
enable	Enables an OSPF accept entry for a specified advertising router. Use the no operator to disable an OSPF accept entry: <code>no accept adv-rtr <A.B.C.D> enable</code>
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. <i>type1</i> means match external type 1 only. <i>type2</i> means match external type 2 only. Use the no operator to disable metric-type: <code>no ip ospf accept adv-rtr <A.B.C.D> metric-type</code>
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 Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya Fabric Connect on VSP Operating System Software*, NN47227-510.

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.

* Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

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>] [vrfs 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 27: 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 28: 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.

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.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.

- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create the redistribution instance:

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

3. 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>]
```

4. Use the following variable definitions table to configure other parameters as required.

5. Enable the redistribution:

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

6. Ensure that the configuration is correct:

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

For RIPng, use `show ipv6 rip redistribute`.

7. 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 29: 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 this document and *Configuring BGP Services on VSP Operating System Software*, NN47227-508.

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 non-contiguous 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 IPv6 Prefix List

Use IPv6 prefix lists to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IPv6**.
2. Click **Policy**.
3. In the **IPv6-Prefix List** tab, click **Insert**.
4. Edit the parameters as required.

5. Click **Insert**.

Ipv6–Prefix list field descriptions

Use the data in the following table to use the **Ipv6–Prefix List** tab.

Name	Description
Id	Specifies the prefix list. The range is 0 to 65535.
Prefix	Specifies the prefix IPv6 address.
PrefixMaskLen	Specifies the length of the prefix mask. You must enter the full 128-bit mask to exact a full match of a specific IPv6 address (for example, when creating a policy to match the next-hop).
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name can be from 1 to 64 characters in length.
MaskLenFrom	Specifies the lower bound on the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Specifies 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 Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya Fabric Connect on VSP Operating System Software*, NN47227-510.

Procedure

1. In the navigation pane, 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	<p>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.</p> <p>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.</p> <p>You can also change this field in the Route Policy tab of the Policy dialog box.</p>
MatchIpRouteDest	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.
MatchNextHop	<p>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.</p> <p>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.</p>
MatchInterface	Specifies if the system matches the IP address of the interface by which the RIP route was learned against the contents of the

Table continues...

Name	Description
	<p>specified prefix list. This field is used only for RIP routes and ignored for all other type of route.</p> <p>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.</p>
MatchRouteType	<p>Configures a specific route type to match (applies only to OSPF routes).</p> <p>Externaltyp1 and Externaltyp2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes. The default is any.</p>
MatchMetric	<p>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.</p>
MatchAsPath	<p>Configures if the system matches the BGP autonomus system path. Applicable to BGP only. This overrides the BGP neighbor filter list information.</p>
MatchCommunity	<p>Filters incoming and outgoing updates based on a Community List. Applicable to BGP only. The default is disable.</p>
MatchCommunityExact	<p>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.</p>
MatchTag	<p>Specifies a list of tags used during the match criteria process. Applicable to BGP only. It contains one or more tag values.</p>
MatchVrf	<p>Identifies the source VRFs that leaks routes to the local VRF (applies only to RIP routes).</p>
NssaPbit	<p>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.</p>
SetRoutePreference	<p>Configures a value from 0 to 255. The default value is 0. If the default is configured, the global preference value is used.</p> <p>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.</p>
SetMetricTypeInternal	<p>Identifies the MED value for routes advertised to BGP numbers to the Interior Gateway Protocol (IGP) metric value. The default is 0.</p>

Table continues...

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	<p>Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged.</p> <ul style="list-style-type: none"> • Unchanged—keeps the community attribute in the route path as it is. • None—removes the community in the route path additive. • Append—adds the community number specified in SetCommunityNumber to the community list attribute.
SetOrigin	Configures the origin for the BGP protocol to IGP, EGP, incomplete, or unchanged. If not configured, the system uses the route origin from the IP routing table (protocol). The default is unchanged.

Table continues...

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 Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya Fabric Connect on VSP Operating System Software*, NN47227-510.

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, 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 pane, 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 effect for an OSPF Redistribute context. This prevents the system from attempting to 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 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.

Table continues...

Name	Description
NextHop	Specifies the IP address of the next hop of this route.
AltSequence	Indicates the alternative route sequence. The value of 0 denotes the best 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	<p>Specifies the router interface for this route.</p> <ul style="list-style-type: none"> • 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	<p>Specifies the routing mechanism through which this route was learned:</p> <ul style="list-style-type: none"> • 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 • bbnSpflgp • 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	<p>Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths.</p> <ul style="list-style-type: none"> • 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

Table continues...

Name	Description
	<ul style="list-style-type: none"> • 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.
NextHopVrflid	Specifies the VRF ID of the next-hop address.

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 Fabric Connect for Layer 3 VSNs and IP Shortcuts, see *Configuring Avaya Fabric Connect on VSP Operating System Software*, NN47227-510.

Procedure

1. In the navigation pane, 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. <ul style="list-style-type: none"> • 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.

*** Note:**

The prompt for the non-PowerPlus chassis is VSP-4850GTS. The prompt for the PowerPlus chassis is VSP-4850GTS-PWR+. The prompt for the Fiber box is VSP-4450 GSX. For consistency, this document uses the VSP-4850GTS prompt.

Enabling routing globally or on a VRF instance

Before you begin

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

About this task

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. Activate IP forwarding:

```
ip routing
```

2. View the forwarding configuration:

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

Example

```
VSP-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1(config)#show ip routing
```

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

```
IP Forwarding is enabled
```

IP routing configuration using ACLI

```
IP ECMP feature is disabled
Maximum ECMP paths number is 1
ECMP 1 pathlist :
ECMP 2 pathlist :
ECMP 3 pathlist :
ECMP 4 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
```

```
VSP-4850GTS-PWR+:1(router-vrf)#show ip routing vrf 1
```

```
=====
                               IP - VRF 1
=====

IP Forwarding is enabled
IP ECMP feature is enabled
Maximum ECMP paths number is 4
ECMP 1 pathlist : path1
ECMP 2 pathlist :
ECMP 3 pathlist :
ECMP 4 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 30: Variable definitions

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

Enabling routing on an IP interface

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 a VLAN or router port.

Procedure

Enable routing:

```
routing enable
```

Example

```
VSP-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1#configure terminal
VSP-4850GTS-PWR+:1 (config)#interface gigabitethernet 1/2
VSP-4850GTS-PWR+:1 (config-if)#routing enable
```

Deleting a dynamically learned route

Before you begin

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

About this task

Delete a dynamically learned route from the routing table if you do not want Virtual Services Platform 4000 to use the route. Exercise caution when you delete entries from the routing table.

Procedure

1. 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<0-16>] [vrfids WORD<0-512>]
[static]
```

2. Delete the dynamically learned route:

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

Example

```
VSP-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1#configure terminal
```

Delete the dynamically learned route:

```
VSP-4850GTS-PWR+:1(config)#no ip route 47.17.10.32 255.255.255.0
47.17.10.31 dynamic
```

Variable definitions

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

Table 31: 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<0-16>	Displays the route for a particular VRF.
vrfids WORD<0-512>	Displays the route for a particular VRF number.
spbm-nh-as-mac	Displays the spbm route next hop as mac.

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

Table 32: 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<0-16>	Specifies the name of the next-hop VRF router.

Configuring IP route preferences

Before you begin

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

- 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

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.

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

Procedure

1. Configure the route preference:

```
ip route preference protocol <static|spbm-level1> <0-255>
```

2. Confirm that the configuration is correct:

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

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Configure the route preference to SPBM Level 1:

```
VSP-4850GTS-PWR+:1(config)#ip route preference protocol spbm-level1 7
```

Confirm the configuration is correct:

```
VSP-4850GTS-PWR+:1(config)# show ip route preference vrf test
```

IP Route Preference - VRF test		
PROTOCOL	DEFAULT	CONFIG
LOCAL	0	0
STATIC	5	5
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.

The default form of the route preference command is `default ip route preference protocol <protocol>`.

Table 33: Variable definitions

Variable	Value
vrf WORD<0-16>	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 GigabitEthernet 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-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+:1#configure terminal
VSP-4850GTS-PWR+:1 (config)#interface gigabitethernet 2/15
VSP-4850GTS-PWR+:1 (config-if)#action flushIp
```

Assigning an IP address to a port

Before you begin

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

About this task

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.

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 <0-127>]
```

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

```
vrf WORD<0-16>
```

3. Confirm that the configuration is correct:

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

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

```
VSP-4850GTS-PWR+:1(config)#interface gigabitethernet 1/11
```

Assign an IP address to the port

```
VSP-4850GTS-PWR+:1(config)#brouter port 1/11 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 34: Variable definitions

Variable	Value
mac-offset <0-127>	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 4000 and is not used if the port is untagged.

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

Table 35: 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 4000 and is not used if the port is untagged.

Assigning an IP address to a VLAN

Before you begin

- Activate IP forwarding globally.
- You must log on to the VLAN Interface Configuration mode in ACLI.

About this task

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

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. Assign an IP address:

```
ip address {A.B.C.D} {A.B.C.D} [<0-127>]
```

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

```
vrf WORD<0-16>
```

Example

```
VSP-4850GTS-PWR+:1>enable
VSP-4850GTS-PWR+1#configure terminal
VSP-4850GTS-PWR+1(config)#interface vlan 2
VSP-4850GTS-PWR+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 36: Variable definitions

Variable	Value
<A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask, respectively.

Table continues...

Variable	Value
<0-127>	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 0–127.

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

Table 37: Variable definitions

Variable	Value
WORD<0-16>	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-4850GTS:1(config)#show ip interface
```

```
=====
                        IP Interface - GlobalRouter
=====
INTERFACE      IP          NET          BCASTADDR  REASM      VLAN  BROUTER
                ADDRESS    MASK          FORMAT      MAXSIZE    ID    PORT
-----
Port1/6        6.6.6.6    255.255.255.0 ones        1500       200   true
Vlan100        5.5.5.5    255.255.255.0 ones        1500       100   false
Vlan4000       47.17.41.21 255.255.255.0 ones        1500       4000  false
=====
```

Variable definitions

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

Table 38: Variable definitions

Variable	Value
gigabitethernet	Displays IP interface information for Gigabit Ethernet ports.
vrf	Displays interface information for a particular VRF.
vrfids	Displays interface information for particular VRF IDs.

Configuring IP routing globally or for a VRF

Before you begin

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

About this task

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.

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

Procedure

1. 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.

2. Activate ECMP:

```
ip ecmp
```

3. Activate the alternative route feature globally:

```
ip alternative-route
```

4. Configure a prefix-list for target destination:

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

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

```
ip ecmp pathlist-<1-4> WORD<1-64>
```

6. Access privileged EXEC mode:

```
end
```

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

```
ip ecmp pathlist-apply
```

8. Configure the remaining global parameters as required.

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Enable ECMP:

```
VSP-4850GTS-PWR+:1(config)#ip ecmp
```

Configure a prefix-list for target destination:

```
VSP-4850GTS-PWR+:1(config)#ip prefix-list LIST1
47.17.121.50/255.255.255.0
```

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

```
VSP-4850GTS-PWR+:1(config)#ip ecmp pathlist-1 LIST1
```

Access privileged EXEC mode:

```
VSP-4850GTS-PWR+:1(config)#end
```

Apply changes to all ECMP path-list apply configurations:

```
VSP-4850GTS-PWR+:1#ip ecmp pathlist-apply
```

Variable definitions

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

Table 39: Variable definitions

Variable	Value
alternative-route	<p>Enables or disables the alternative route feature. The default value is enabled.</p> <p>If the alternative-route parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are readded.</p> <p>The default form of this command is <code>default ip alternative-route</code>. The no form of this command is <code>no ip alternative-route</code>.</p>
max-routes-trap enable	<p>Enables the device to send a trap after the maximum number of routes is exceeded.</p> <p>The no form of this command is <code>no max-routes-trap enable</code>. The default form of this command is <code>default max-routes-trap enable</code>.</p>
more-specific-non-local-route	<p>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.</p>

Table continues...

Variable	Value
	The default form of this command is <code>default ip more-specific-non-local-route</code> . The no form of this command is <code>no ip more-specific-non-local-route</code> .
routing	Enables routing. The no form of this command is <code>no ip routing</code> .
supernet	Enables or disables supernetting. If you globally enable supernetting, the device can learn routes with a route mask of less than 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. The default form of this command is <code>default ip supernet</code> . The no form of this command is <code>no ip supernet</code> .
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 <code>default ip ttl</code> .

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

Table 40: Variable definitions

Variable	Value
pathlist-1 <i>WORD</i> <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 <code>ip ecmp</code> . The no form of this command is <code>no ip ecmp pathlist-1</code> .
pathlist-2 <i>WOR</i> <0-64> <i>WORD</i> <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 <code>ip ecmp</code> . The no form of this command is <code>no ip ecmp pathlist-2</code> .
pathlist-3 <i>WORD</i> <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 <code>ip ecmp</code> . The no form of this command is <code>no ip ecmp pathlist-3</code> .
pathlist-4 <i>WORD</i> <0-64>	Configures up to four equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.

Table continues...

Variable	Value
	To configure this parameter, you must globally activate ECMP, with the command <code>ip ecmp</code> . The no form of this command is <code>no ip ecmp pathlist-4</code> .
max-path <1-4>	Configures the maximum number of ECMP paths. The range for this number 1–4. The default form of this command is <code>default ip ecmp max-path</code> .

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

Table 41: Variable definitions

Variable	Value
redirect	Enables the device to send ICMP destination redirect messages. The default form of this command is <code>default ip icmp redirect</code> . This setting is disabled by default.
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 <code>default ip icmp unreachable</code> .

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
```

- Use the following variable definitions table to configure other static route parameters as required.

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

```
show ip route static
```

- Delete a static route:

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

The limit for a static route for VSP 4000 is 1000.

Example

```
VSP-4850GTS-PWR+:1>enable
```

Log on to Global Configuration mode:

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create an IP static route:

```
VSP-4850GTS-PWR+:1(config)#ip route 42.17.0.0 255.255.0.0 42.17.156.126 weight 200
```

Enable a static route:

```
VSP-4850GTS-PWR+: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-4850GTS-PWR+:1(config)#show ip route static
```

Variable definitions

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

Table 42: 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.
disable	Disables a route to the router or VRF.
enable	Adds a static or default route to the router or VRF.

Table continues...

Variable	Value
	<p>The no form of this command is <code>no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> enable</code>.</p> <p>The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> enable</code>.</p>
local-next-hop enable	<p>Enables the local next hop for this static route. The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> local-next-hop enable</code>.</p> <p>The no form of this command is <code>no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> local-next-hop enable</code>.</p>
next-hop-vrf <WORD 0-16>	<p>Specifies the next-hop VRF instance by name.</p> <p>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).</p> <p>The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> next-hop-vrf <WORD 0-16></code>.</p> <p>The no form of this command is <code>no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> next-hop-vrf <WORD 0-16></code>.</p>
weight <1-65535>	<p>Specifies the static route cost.</p> <p>The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> weight</code>.</p>
preference <1-255>	<p>Specifies the route preference.</p> <p>The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> preference</code>.</p>

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

Table 43: 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<0-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configuring a black hole static route

Before you begin

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

About this task

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. Create a black hole static route:

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

2. 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<0-16>]
```

3. 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-4850GTS-PWR+:1>enable
```

Log on to Global Configuration mode:

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create a black hole static route:

```
VSP-4850GTS-PWR+: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-4850GTS-PWR+: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 44: 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 <code>no ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 enable.</code>
local-next-hop enable	Enables the local next hop for this static route. The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> local-next-hop enable.</code> The no form of this command is <code>no ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 local-next-hop enable.</code>
next-hop-vrf WORD<0-16>	Specifies the next-hop VRF instance by name. The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 next-hop-vrf <WORD 0-16>.</code> The no form of this command is <code>no ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 next-hop-vrf <WORD 0-16>.</code>
weight <1-65535>	Specifies the static route cost. The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 weight.</code>
preference <1-255>	Specifies the route preference. The default form of this command is <code>default ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 preference.</code>

Configuring a default static route

Before you begin

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

About this task

The default route specifies 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 4000 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. Create a default static route:

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

2. Enable a default static route:

```
ip route 0.0.0.0 0.0.0.0 <A.B.C.D> enable [next-hop-vrf WORD<0-16>]
```

3. Configure other default static route parameters as required.

Example

```
VSP-4850GTS-PWR+:1>enable
```

Log on to Global Configuration mode:

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create a default static route:

```
VSP-4850GTS-PWR+:1(config)#ip route 0.0.0.0 0.0.0.0 42.17.159.128 weight 100
```

Enable a default static route:

```
VSP-4850GTS-PWR+: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 45: 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 <code>no ip route 0.0.0.0 0.0.0.0 <A.B.C.D> enable</code> .
local-next-hop enable	Enables the local next hop for this static route. The default form of this command is <code>default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> local-next-hop enable</code> . The no form of this command is <code>no ip route 0.0.0.0 0.0.0.0 <A.B.C.D> local-next-hop enable</code> .
next-hop-vrf WORD<0-16>	Specifies the next-hop VRF instance by name. The default form of this command is <code>default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> next-hop-vrf WORD<0-16></code> .

Table continues...

Variable	Value
	The no form of this command is <code>no ip route 0.0.0.0 0.0.0.0 <A.B.C.D> next-hop-vrf WORD<0-16></code> .
weight <1-65535>	Specifies the static route cost. The default form of this command is <code>default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> weight</code> .
preference <1-255>	Specifies the route preference. The default form of this command is <code>default ip route 0.0.0.0 0.0.0.0 <A.B.C.D> preference</code> .

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<0-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-4850GTS-PWR+:1>enable
```

Log on to Global Configuration mode:

```
VSP-4850GTS-PWR+:1#configure terminal
```

Enable ICMP Router Discovery of the device:

```
VSP-4850GTS-PWR+:1(config)#ip irdp
```

confirm that Router Discovery is enabled:

```
VSP-4850GTS-PWR+:1(config)#show ip irdp
```

Variable definitions

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

Table 46: Variable definitions

Variable	Value
interface	Displays route discovery interface information.
vrf <i>WORD</i> <0–16>	Displays route discovery for particular VRF.
vrfids <i>WORD</i> <0–512>	Displays route discovery for particular VRF IDs.

Enabling or disabling IPv4 ICMP broadcast globally

On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when they reach the control plane.

About this task

Use these commands to enable or disable the IPv4 ICMP broadcast feature on the global router.

Procedure

1. Enter Global Configuration mode:

```
enable  
configure terminal
```
 2. Enable IPv4 ICMP broadcast feature, enter:

```
ip icmp echo-broadcast-request
```
 3. Disable IPv4 ICMP broadcast feature, enter:

```
no ip icmp echo-broadcast-request
```
 4. Set the IPv4 ICMP broadcast feature to default state, enter:

```
default ip icmp echo-broadcast-request
```
- * Note:**
By default, the IPv4 ICMP broadcast feature is enabled.
5. View the IPv4 ICMP broadcast feature state:

```
show ip routing
```

Enabling or disabling IPv4 ICMP broadcast per VRF

On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when they reach the control plane.

About this task

Use these commands to enable or disable the IPv4 ICMP broadcast feature on the VRF router.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Enable IPv4 ICMP broadcast feature, enter:

```
ip icmp echo-broadcast-request
```

3. Disable IPv4 ICMP broadcast feature, enter:

```
no ip icmp echo-broadcast-request
```

4. Set the IPv4 ICMP broadcast feature to default state, enter:

```
default ip icmp echo-broadcast-request
```

 **Note:**

By default, the IPv4 ICMP broadcast feature is enabled.

5. View the IPv4 ICMP broadcast feature state:

```
show ip routing
```

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-4850GTS-PWR+:1>enable
```

Log on to Global Configuration mode:

```
VSP-4850GTS-PWR+:1#configure terminal
```

Log on to the GigabitEthernet Interface mode:

```
VSP-4850GTS-PWR+:1(config)#interface gigabitethernet 1/16
```

Specify the address placed in advertisement packets to the all-systems multicast address:

```
VSP-4850GTS-PWR+:1(config-if)#ip irdp address 244.0.0.1
```

Enable the interface to send the advertisement packets:

```
VSP-4850GTS-PWR+:1(config-if)#ip irdp multicast
```

Configure the lifetime for advertisements:

```
VSP-4850GTS-PWR+:1(config-if)#ip irdp holdtime 180
```

Variable definitions

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

Table 47: 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 <code>default ip irdp address</code> .
holdtime <4-9000>	Configures the lifetime for advertisements. The default form of this command is <code>default ip irdp holdtime</code> .
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 <code>default ip irdp maxadvertinterval</code> .

Table continues...

Variable	Value
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 <code>default ip irdp minadvertinterval</code> .
multicast	Specifies if multicast advertisements are sent. The no form of this command is <code>no ip irdp multicast</code> .
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 <code>default ip irdp preference</code> .

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.

* Note:

You can configure only one CLIP interface with an IPv6 address, which can be only used as a source IPV6 address for IS-IS.

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<0-16>]
```

3. View the IP address on the CLIP interface:

```
show ip interface
```

Example

```
VSP-4850GTS-PWR+:1>enable
```

Log on to Global Configuration mode:

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create or access a CLIP interface:

```
VSP-4850GTS-PWR+:1(config)#interface loopback 200
```

Configure an IP address for the interface:

```
VSP-4850GTS-PWR+:1(config-if)#ip address 200 45.17.159.120/255.255.0.0
```

```
VSP-4850GTS-PWR+:1(config-if)#show ip interface
```

Variable definitions

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

Table 48: Variable definitions

Variable	Value
address [<1-256>] <A.B.C.D/X> [vrf WORD<0-16>]	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<0-16> specifies an associated VRF by name. The no form of this command is no ip address [<1-32>] <A.B.C.D> [vrf WORD<0-16>].
area <1-256> <A.B.C.D>[vrf WORD<0-16>]	Designates an area for the CLIP interface. vrf WORD<0-16> specifies an associated VRF by name The default form of this command is default ip area <1-256> <A.B.C.D> [vrf WORD<0-16>]. The no form of this command is no ip area <1-256> vrf WORD<0-16>].
ospf [<1-256>] [vrf <WORD 0-16>]	Enables OSPF for the CLIP interface. <1-256> specifies the interface. vrf <WORD 0-16> specifies an associated VRF by name. The default form of this command is default ip ospf <1-256> [vrf <WORD 0-16>]. The no form of this command is no ip ospf <1-256> [vrf <WORD 0-16>].

Table continues...

Variable	Value
pim [<1-256>] [bsr-candidate preference <0-255>	<p>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.</p> <p>The default form of this command is <code>default ip pim <1-256> [bsr-candidate]</code>. The no form of this command is <code>no ip pim <1-256> [bsrcandidate]</code>.</p>

Creating an IPv6 CLIP interface

About this task

Create an IPv6 CLIP interface

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Specify an interface ID value:


```
interface loopback <1-256>
```
3. Create an IPv6 loopback interface address:


```
ipv6 interface address WORD <0-255>
```
4. Ensure the configuration is correct:


```
show ipv6 interface loopback
```

Example

```
Switch:1#show ipv6 interface loopback
```

```

=====
                        Loopback IPv6 Interface
=====
IF   Descr          VLAN  PHYSICAL          ADMIN  OPER  TYPE  MTU  HOP  REACHABLE  RETRANSMIT
INDX                                     STATE  STATE                                     LMT  TIME      TIME
-----
1344 CLIPv6-1      --    00:00:00:00:00:01 enable  up    ETHER 1500  64   30000      1000
=====
                        Loopback IPv6 Address
=====
IPV6 ADDRESS/PREFIX LENGTH          LOOPBACK-ID  TYPE    ORIGIN  STATUS
-----
1:210:0:0:0:0:0:210/128              C-1          UNICAST MANUAL  PREFERRED

```

Legend: NA - Information not available

1 out of 204 Total Num of Interface Entries displayed.
1 out of 407 Total Num of Address Entries displayed.

Variable definitions

Use the data in the following table to use the `ipv6` commands.

Variable	Value
WORD<1–256>	Specifies the CLIP interface ID.
WORD<0–255>	Specifies the IPv6 address.

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**.

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.
AltSequence	Indicates the alternative route sequence. The value of 0 denotes the best route.
NextHopId	Specifies the identifier of the next-hop, hostname or MAC address.
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. <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • local—nonprotocol information, for example, manually configured entries • static • isis • inter-vrf redistributed route
Age	Specifies the number of seconds since this route was last updated or otherwise determined correct.

Table continues...

Name	Description
PathType	<p>Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths.</p> <ul style="list-style-type: none"> • 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 • 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.
NextHopVrflid	Specifies the VRF ID of the next-hop address.

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 **ConfiguredValue** 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
DefaultValue	Specifies the default preference value for the specified protocol.
Protocol	Specifies the protocol name.
ConfiguredValue	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. <ul style="list-style-type: none"> • 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.

Table continues...

Name	Description
MacOffset	Specifies a number by which to offset the MAC address of the VLAN from the chassis MAC address. This ensures that each IP address has a different MAC address.
Vrflid	Specifies the associated VRF interface. The Vrflid 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 237.

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 > Basic**.
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. <ul style="list-style-type: none"> • 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 4000 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 VLAN 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 switch 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 IPv4 ICMP broadcast, select **IcmpEchoBroadcastRequestEnable**.
7. To globally enable the Alternative Route feature, select **AlternativeEnable**.
8. To globally enable ICMP Router Discovery, select **RouteDiscoveryEnable**.
9. To globally enable ECMP, select **EcmpEnable**.
10. Configure the remaining parameters as required.
11. 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	<p>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.</p> <p> Important:</p> <p>Avaya recommends that you only enable icmp-unreach-msg if it is absolutely required. If icmp-unreach-msg is enabled</p>

Table continues...

Name	Description
	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. The default is enabled.
IcmpEchoBroadcastRequestEnable	Enables or disables IP ICMP echo broadcast request feature. The default is enabled.
AlternativeEnable	Globally enables or disables the Alternative Route feature. 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 4000. The default value is 360 minutes. The range for this value is 1 to 32767 minutes.
ArpThreshold	Defines the maximum number of outstanding ARP requests a device can generate. The default is 500.
ArpMcastMacFlooding	Enables or disables IP ARP multicast MAC flooding. The default is disabled.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled.

Table continues...

Name	Description
	After ECMP is disabled, the <code>EcmpMaxPath</code> is reset to the default value of 1.
EcmpMaxPath	<p>Globally configures the maximum number of ECMP paths.</p> <ul style="list-style-type: none"> The interval is 1 to 4. The default value is: <ul style="list-style-type: none"> - 1 when ECMP is disabled - 4 when ECMP is enabled <p>You cannot configure this feature unless ECMP is enabled globally.</p>
Ecmp1PathList	Selects a preconfigured ECMP path.
Ecmp2PathList	Selects a preconfigured ECMP path.
Ecmp3PathList	Selects a preconfigured ECMP path.
Ecmp4PathList	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 four 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 237.

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 1000.

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 **OwnerVrflid** 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 **NextHopVrflid** 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
OwnerVrflid	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.
NextHopVrflid	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.

Table continues...

Name	Description
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 237.

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 4000 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 **OwnerVrflid** 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 237.

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 **OwnerVrflid** 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 237.

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.
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.
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 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 237.

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.

Table continues...

Name	Description
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 237.

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.
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 Circuitless IPv4 interface

About this task

You can use a circuitless IPv4 (CLIPv4) interface to provide uninterrupted connectivity to your system.

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.
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.

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.

Table continues...

Name	Description
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.

Table continues...

Name	Description
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: <ul style="list-style-type: none"> • 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 pane, expand the following folders: **Configuration > IP** or **Configuration > IPv6**.
2. 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

Perform this procedure to configure RSMLT on each IP VLAN interface.

Before you begin

- 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

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

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4059>
```

2. Enable RSMLT on a VLAN:

```
ip rsmlt
```

Example

```
Switch:1>enable
```

```
Switch:1#configure terminal
```

```
Log on to VLAN Interface Configuration mode:
```

```
Switch:1(config)#interface VLAN 100
```

```
Enable RSMLT on a VLAN:
```

```
Switch:1(config-if)#ip rsmlt
```

Variable definitions

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

Table 49: Variable definitions

Variable	Value
holddown-timer <0-3600>	<p>Configures how long the RSMLT device does not participate in Layer 3 forwarding.</p> <p><i>0-3600</i> is the timer value in seconds.</p> <p>To configure this value to the default value, use the default operator with this command.</p> <p>Configure this value to be longer than the anticipated routing protocol convergence.</p>
holdup-timer <0-3600 9999>	<p>Configures how long the RSMLT device maintains forwarding for its peer.</p> <p><i>0-3600 9999</i> is the timer value in seconds. 9999 means infinity.</p> <p>To configure this value to the default value, use the default operator with this command.</p>

Showing IP RSMLT information

Show IP RSMLT information to view data about all RSMLT interfaces.

About this task

Important:

If you use the `show ip rsmlt` command after you delete an RSMLT, the RSMLT still shows until you restart the switch.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display RSMLT information using the following command:

```
show ip rsmlt {edge-support} [<local|peer>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```


Example

```
Switch:1>enable
Switch:1#show ip rsmlt
```

```
=====
                          Ip Rsmlt Local Info - GlobalRouter
=====
VID      IP                MAC                ADMIN  OPER  HDTMR  HUTMR
-----
1000    100.0.0.12          b0:ad:aa:40:05:25  Enable Up    60     180
1500    150.0.0.12          b0:ad:aa:40:05:28  Enable Up    60     180
3000    200.0.0.12          b0:ad:aa:40:05:01  Enable Up    60     180

VID      SMLT ID
-----
1000     50
1500     50
3000

VID      IPv6                MAC                ADMIN  OPER  HDTMR  HUTMR
-----
1000                                b0:ad:aa:40:05:25  Enable Up    60     180
      100:0:0:0:0:0:0/64
      100:0:0:0:0:0:0:12/64
      fe80:0:0:0:b2ad:aaff:fe40:525/128
1500                                b0:ad:aa:40:05:28  Enable Up    60     180
      150:0:0:0:0:0:0/64
      150:0:0:0:0:0:0:12/64
      fe80:0:0:0:b2ad:aaff:fe40:528/128
3000                                b0:ad:aa:40:05:01  Enable Up    60     180
      30:0:0:0:0:0:0/64
      30:0:0:0:0:0:0:12/64
      fe80:0:0:0:b2ad:aaff:fe40:501/128

VID      SMLT ID
-----
1000     50
1500     50
3000

=====
                          Ip Rsmlt Peer Info - GlobalRouter
=====
VID      IP                MAC                ADMIN  OPER  HDTMR  HUTMR
-----
1000    100.0.0.206         b0:ad:aa:41:7d:23  Enable Up    60     180
1500    150.0.0.206         b0:ad:aa:41:7d:24  Enable Up    60     180

VID      HDT REMAIN  HUT REMAIN  SMLT ID
-----
1000     60          180         50
1500     60          180         50

VID      IPv6                MAC                ADMIN  OPER  HDTMR  HUTMR
-----
1000                                b0:ad:aa:41:7d:23  Enable Up    60     180
      100:0:0:0:0:0:0/64
      100:0:0:0:0:0:0:206/64
      fe80:0:0:0:b2ad:aaff:fe41:7d23/128
```

```

1500          b0:ad:aa:41:7d:24  Enable  Up    60    180
150:0:0:0:0:0:0/64
150:0:0:0:0:0:206/64
fe80:0:0:0:b2ad:aaff:fe41:7d24/128

VID  HDT  REMAIN  HUT  REMAIN  SMLT  ID
-----
1000  60          180    50
1500  60          180    50

Switch:1#

```

Variable definitions

Use the information in the following command to use the `show ip rsmlt` command.

Table 50: 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 51: Variable definitions

Variable	Value
VID	Indicates the VLAN ID.
IP	Indicates the IP address of the VLAN.
MAC	Indicates the MAC address assigned.
ADMIN	Indicates the administrative status of RSMLT on the VLAN.
OPER	Indicates the operational status of RSMLT on the VLAN.
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.
SMLT ID	Indicates the Split MultiLink Trunk ID.

Configuring RSMLT edge support

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.

About this task

RSMLT edge support is disabled by default.

Important:

If you use the `show ip rsmlt` command after you delete an RSMLT, the RSMLT still displays until you restart the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable RSMLT-edge:

```
ip rsmlt edge-support
```

Use the `no` operator to disable RSMLT-edge: `no ip rsmlt edge-support`

3. Clear RSMLT peer information, and then delete the RSMLT peer address:

```
no ip rsmlt peer-address <1-4059>
```

4. Display RSMLT-edge status information:

```
show ip rsmlt edge-support
```

Example

```
Switch:1>enable
```

```
Switch:1#configure terminal
```

Enable RSMLT-edge:

```
Switch:1(config)#ip rsmlt edge-support
```

Display RSMLT-edge status information:

```
Switch: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.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. VLAN IDs 1 to 4059 are configurable. The system reserves VLAN IDs 4060 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

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.

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

Use the following procedure to configure RSMLT using EDM

Procedure

1. In the navigation pane, 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	<p>Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address.</p> <p>The range of this value is from 0 to 3600 seconds. The default is 60.</p> <p>If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.</p>
HoldUpTimer	<p>Defines how long the RSMLT system maintains forwarding for its peer if the peer is down. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 180.</p> <p>If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.</p>

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	IP interface identification.
VlanId	Specifies the VLAN ID of the chosen VLAN.

Table continues...

Name	Description
IpAddr	Specifies the IP address on the RSMLT VLAN.
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	IP interface identification.
VlanId	Specifies the VLAN ID of the chosen VLAN.
IpAddr	Specifies the IP address on the RSMLT VLAN.
MacAddr	Specifies the MAC address of the selected VLAN.
Enable	Displays the RSMLT operating status as enabled or disabled.

Table continues...

Name	Description
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.
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.

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.
PeerVrflid	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.

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:**

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

! Important:

The switch, 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 the switch 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.

*** Note:**

The VSP VRRP IP address responds only to ICMP-based traceroute requests. It does not respond to UDP-based traceroute requests.

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.

Configuring VRRP on a port or a VLAN

About this task

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.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

*** Note:**

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure a backup VRRP address:

```
ip vrrp address <1-255> <A.B.C.D>
```

3. Configure VRRP on a port:

```
ip vrrp <1-255> enable
```

4. Show the global VRRP configuration:

```
show ip vrrp
```

Example

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

```
Switch:1(config)# interface gigabitethernet 1/2
```

Configure a backup VRRP address:

```
Switch:1(config-if)# ip vrrp address 28 45.16.17.2
```

Configure VRRP on a port:

```
Switch:1(config-if)# ip vrrp 28 enable
```

Show the global VRRP configuration:

```
Switch:1(config-if)# show ip vrrp
```

Variable definitions

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

Table 52: Variable definitions



Variable	Value
1-255	Specifies the number of the VRRP to create or modify.
action {none preempt}	<p>Causes the virtual router to disregard the timer and transition to Master state immediately, provided the hold-down timer is running.</p> <p> Note: You can use this parameter only if the hold-down timer is active.</p> <p>To set this option to the default value, use the default operator with this command.</p>
address <1-255> <A.B.C.D>	<p>Configures the IP address of the VRRP physical interface that forwards packets to the virtual IP addresses associated with the virtual router.</p> <p>A.B.C.D is the IP address of the master VRRP.</p> <p>Use the no operator to remove the IP address of the VRRP physical interface: <code>no ip vrrp address <1-255> <A.B.C.D></code>. To configure this option to the default value, use the default operator with this command.</p>
adver-int <1-255>	<p>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.</p> <p>To configure this option to the default value, use the default operator with this command.</p>
backup-master enable	<p>Enables the VRRP backup master.</p> <p>Use the no operator to disable the VRRP backup master: <code>no ip vrrp <1-255> backup-master enable</code>. To configure this option to the default value, use the default operator with this command.</p> <p>When backup master functionality is enabled, the VRRP router will IP-forward packets destined to the VRRP MAC even when the router is not the VRRP Master.</p> <p> Important: Do not enable backup master if you enable critical IP.</p>
critical-ip-addr <A.B.C.D>	<p>Configures the critical IP address for VRRP.</p> <p>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 switch in the virtual router (for example, from master to backup) in case the interface stops responding.</p>

Table continues...



Variable	Value
	<p> Note:</p> <p>In this context, <i>local</i> implies an address from the same VRF as the IP interface where VRRP is being configured.</p>
critical-ip enable	<p>Enables the critical IP address option.</p> <p>Use the no operator to disable the critical IP address option: <code>no ip vrrp <1-255> critical-ip enable</code>. To configure this option to the default value, use the default operator with this command.</p> <p> Important:</p> <p>Do not enable Critical IP if backup master is enabled.</p>
enable	<p>Enables VRRP on the port.</p> <p>Use the no operator to disable VRRP on the port: <code>no ip vrrp <1-255> enable</code>. To configure this option to the default value, use the default operator with this command.</p>
fast-adv enable	<p>Enables the Fast Advertisement Interval. The default is disabled.</p> <p>Use the no operator to disable VRRP on the port: <code>no ip vrrp <1-255> fast-adv enable</code>. To configure this option to the default value, use the default operator with this command.</p>
fast-adv-int <200-1000>	<p>Configures the Fast Advertisement Interval, the time interval between sending VRRP advertisement messages.</p> <p>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.</p> <p>To configure this option to the default value, use the default operator with this command.</p>
holddown-timer <0-21600>	<p>Specifies the time interval (in seconds) for which the transition of virtual router to Master state is delayed in case of the following conditions:</p> <ul style="list-style-type: none"> • The VRRP hold-down timer runs only when the VRRP virtual router transitions from initialization to backup to master. This occurs only on a system startup. • The VRRP hold-down timer does not run if the amount of time passed since VRRP virtual router initialization is greater than preset hold-down time. In such a case, VRRP virtual router transitions to Master happens irrespective of the hold-down timer. • The VRRP hold-down timer also applies to the VRRP BackupMaster feature. <p>0-21600 is the time interval range (in seconds). To configure this option to the default value, use the default operator with this command. The default value for hold-down timer is 0, that is, the timer is disabled by default.</p>

Table continues...

Variable	Value
priority <1-255>	<p>Configures the port VRRP priority.</p> <p>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.</p> <p>To configure this option to the default value, use the default operator with this command.</p>

Showing VRRP information

About this task

Show VRRP port or VLAN information to view configuration details and operational status.

Procedure

1. Enter Privileged EXEC mode:
enable
2. 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>]
3. Displaying the VRRPv3 configuration:
show ip vrrp address version <2-3>
4. Displaying version based VRRP configuration for the specified VRF:
show ip vrrp address vrf WORD<1-16> version <2-3>
5. Displaying version based VRRP configuration for the specified VRF ID:
show ip vrrp address vrfids WORD<0-512> version <2-3>

Example

```
Switch:1#show ip vrrp address
=====
VRRP Info - GlobalRouter
=====
VRRP ID  P/V      IP           MAC           STATE    CONTROL  PRIO  ADV VERSION
-----
3         3         30.30.30.99  00:00:5e:00:01:03  Master   Enabled  100   1   2
2         1/1      20.20.20.99  00:00:5e:00:01:02  Master   Enabled  100   1   3
2 out of 2 Total Num of VRRP Address Entries displayed.
VRRP ID  P/V      MASTER      UP TIME      HLD DWN  CRITICAL IP(ENABLED)  VERSION
```

```

-----
3      3      30.30.30.18  0 day(s), 00:08:53  0      0.0.0.0      (No)  2
2      1/1     20.20.20.18  0 day(s), 00:02:01  0      0.0.0.0      (No)  3

2 out of 2 Total Num of VRRP Address Entries displayed.

VRRP ID  P/V      BACKUP MASTER  BACKUP MASTER STATE  FAST ADV (ENABLED)  VERSION
-----
3      3      disable        down                  200      (NO)            2
2      1/1     disable        down                  200      (NO)            3

2 out of 2 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 53: 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.
version <2–3>	Specifies the VRRP version (2 or 3) to be shown.

Use the data in the following table to interpret the `show ip vrrp address` command output.

Table 54: Field descriptions

Name	Description
ADV	Indicates the Advertisement Interval, in seconds, between sending advertisement messages.
BACKUP MASTER	Indicates if the Backup-Master feature is disabled or enabled.
BACKUP MASTER STATE	Indicates if the Backup-Master is up. If the switch is in Master state but Backup-Master is enabled, then the BACKUP MASTER STATE will be down.
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.

Table continues...

Name	Description
CRITICAL IP	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.
CRITICAL IP (ENABLED)	Indicates if the critical IP feature is enabled.
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	<p>Specifies the time interval (in seconds) the Hold-down timer has until it expires. If the value is 0, it means the Hold-down timer is not running. This timer will delay the transition from Backup to Master only on a system startup (the VRRP comes from INIT to Backup and determines it should become Master).</p> <ul style="list-style-type: none"> • 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
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.
PRIO	<p>Indicates the priority for the virtual router with respect to other virtual routers that are backing up one or more associated IP addresses. Higher values indicate higher priority.</p> <p>A priority of 255 cannot be configured and it is set for the VRRP router that has the same IP as the physical IP addresses (is Address Owner).</p>
P/V	Indicates the P(ort)/V(lan) on which the VRRP was configured.
STATE	<p>Indicates the current state of the virtual router.</p> <p>initialize—waiting for a startup event</p> <p>backup—monitoring the state or availability of the master router</p> <p>master—forwarding IP addresses associated with this virtual router.</p>
UP TIME	Indicates the time interval since this virtual router exited the INIT state.
VRID	Indicates the virtual router ID on a VRRP router.

Table continues...

Name	Description
VERSION	Indicates the VRRP version.

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 [<1-4059>] [portList] verbose [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1#show ip vrrp interface vlan
```

```
=====
                                Vlan Vrrp
=====
VLAN VRF          VRRP IP          VIRTUAL
ID  NAME          ID   ADDRESS          MAC ADDRESS
-----
200 GlobalRouter  17   9.9.9.42         00:00:5e:00:01:11

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.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. VLAN IDs 1 to 4059 are configurable. The system reserves VLAN IDs 4060 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
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 [<1-4059>] [portList] verbose [vrf WORD<1-16>] [vrfids WORD<0-512>]` command output.

Variable	Value
VLAN ID	Indicates the VLAN ID.
STATE	Indicates the current state of the virtual router. <ul style="list-style-type: none"> • 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
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 VRRPv3 information about the specified interface:
show ip vrrp interface version <2-3>
3. Display additional VRRPv3 information about the specified interface:
show ip vrrp interface verbose version <2-3>
4. Display VRRPv3 information for the specified VRF:
show ip vrrp interface vrf WORD<1-16> version <2-3>
5. Display VRRPv3 information for the specified virtual router:
show ip vrrp interface vrfids WORD<0-512> [version <2-3>]

Example

```
Switch:1#show ip vrrp interface
=====
                        Vlan Vrrp
=====
VLAN VRF          VRRP IP          VIRTUAL          VERSION
ID  NAME          ID  ADDRESS        MAC ADDRESS
-----
3   GlobalRouter  3   30.30.30.99    00:00:5e:00:01:03 2
All 1 out of 1 Total Num of Vlan Vrrp displayed
=====

                        Port Vrrp
=====
PORT VRF          VRRP IP          VIRTUAL          VERSION
NUM  NAME          ID  ADDRESS        MAC ADDRESS
-----
1/1  GlobalRouter  2   20.20.20.99    00:00:5e:00:01:02 3
Switch:1#

Switch:1#show ip vrrp interface verbose
=====
                        Vlan Vrrp Extended
=====
VLAN VRRP VRF          STATE  CONTROL PRIORITY  MASTER  ADVERTISE CRITICAL  VERSION
ID  ID  NAME          STATE  CONTROL PRIORITY  IPADDR  INTERVAL  IPADDR
-----
10  1   Global~  init  disable 100    0.0.0.0    1    0.0.0.0    3
20  2   Global~  init  disable 100    0.0.0.0    1    0.0.0.0    3
All 2 out of 2 Total Num of Vlan Vrrp Extended Entries displayed

VLAN VRRP VRF          HOLDDWN ACTION  CRITICAL BACKUP  BACKUP  FAST ADV  FAST ADV VERSION
ID  ID  NAME          TIME   TIME   IP      MASTER  MASTER  INTERVAL  ENABLE  VERSION
```

```

-----
                                ENABLE          STATE
-----
10  1  GlobalRouter 0      none  disable  disable  down   200    disable  3
20  2  GlobalRouter 0      none  disable  disable  down   200    disable  3

All 2 out of 2 Vlan Vrrp Extended Entries displayed

VLAN VRRP VRF          MASTER ADV  PREEMPT  PSEUDO-HEADER VERSION
ID  ID  NAME          INTERVAL(ms)  MODE      CHECKSUM
-----
10  1  GlobalRouter 1000          enabled  enabled   3
20  2  GlobalRouter 1000          enabled  enabled   3

All 2 out of 2 Vlan Vrrp Extended Entries displayed

=====
                                Port Vrrp Extended
=====
=====
PORT  VRRP VRF          MASTER          ADVERTISE CRITICAL          VERSION
NUM  ID  NAME      STATE  CONTROL  PRIORITY  IPADDR          INTERVAL  IPADDR
-----
1/2  3  Global~  init   disable 100        0.0.0.0         1          0.0.0.0         3

PORT  VRRP VRF          HOLDDWN ACTION  CRITICAL BACKUP  BACKUP  FAST ADV  FAST ADV VERSION
NUM  ID  NAME          TIME      IP      MASTER  MASTER  INTERVAL  ENABLE
          ENABLE          STATE
-----
1/2  3  GlobalRouter 0      none  disable  disable  down   200    disable  3

PORT  VRRP VRF          MASTER ADV  PREEMPT  PSEUDO-HEADER VERSION
NUM  ID  NAME      INTERVAL(ms)  MODE      CHECKSUM
-----
1/2  3  GlobalRouter 1000          enabled  enabled   3

```

Variable definitions

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

Table 55: Variable definitions

Variable	Value
<code>gigabitethernet {slot/port[-slot/port][,...]}</code>	Specifies to show the VRRP information of which interface.
<code>verbose</code>	Specifies to show all available information about the VRRP interfaces.
<code>vlan</code>	Specifies the VLAN that contains the VRRP.
<code>vrf WORD<1-16></code>	Specifies the name of the VRF.
<code>vrid <1-255></code>	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.
<code>vrfids WORD<0-512></code>	Specifies the ID of the VRF and is an integer in the range of 0–512.
<code>version<2–3></code>	Specifies the VRRP version (2 or 3) configured.

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<1-16>]
default ping-virtual-address enable [vrf WORD<1-16>]
```
3. Disable ping to a virtual IP address:


```
no ping-virtual-address enable [vrf WORD<1-16>]
```
4. Display the configuration:


```
show ip vrrp [vrf WORD<1-16>]
```

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<1-16>	Specifies the VRF.

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 of Avaya Virtual Services Platform 4000 Series*, NN46251-700

Procedure

1. Enter VRRP Router Configuration mode:


```
enable
configure terminal
router vrrp
```
2. Enable a trap for VRRP events:


```
send-trap enable [vrf WORD<1-16>]
```
3. Disable a trap for VRRP events:


```
no send-trap enable [vrf WORD<1-16>]
```
4. Configure a trap for VRRP events to the default:


```
default send-trap enable [vrf WORD<1-16>]
```
5. Display the configuration:


```
show ip vrrp [vrf WORD<1-16>]
```

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 <i>WORD</i> <1–16>	Configures the send-trap for a particular VRF.

Configuring VRRP version on an interface

About this task

Use the following command to configure the VRRP version on an interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-
port]][, ...]} OR interface vlan <1-4059>
```

Note:

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Use the following command to configure the VRRP version:

```
ip vrrp version <2-3>
```

3. Use the following command to set the VRRP version to default:

```
default ip vrrp version
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Configure VRRP version for the specified interface:

```
Switch:1(config-if)# ip vrrp version 3
```

Variable definitions

Use the data in the following table to use the `ip vrrp version` command.

Variable	Value
<i>version</i> <2–3>	Configures the VRRP version (2 or 3) on the specified interface

Enabling IPv4 VRRP preempt-mode

You can configure VRRP to preempt the existing router. If a new VRRP router is added to the network with a higher priority than the existing routers, then the new router becomes the master. If preempt-mode is disabled, then the new router does not become a master, it transitions to master only when the current master is down, that is when it does not receive any advertisement packets from the current master. By default, preempt-mode is enabled.

Procedure

1. Enter Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-
port]][, ...]} OR interface vlan <1-4059>
```

*** Note:**

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enter the following command:

```
ip vrrp <vrid> preempt-mode enable
```

3. Use the following command to set the preempt-mode to its default value:

```
default ip vrrp <vrid> preempt-mode
```

4. Use the following command to disable the preempt-mode:

```
no ip vrrp <vrid> preempt-mode enable
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Enabling preempt-mode on interface 1/2:

```
Switch:1(config-if)# ip vrrp 1 preempt-mode enable
```

Variable definitions

Use the data in the following table to use the `ip vrrp <vrid>` command.

Variable	Value
<code>preempt-mode enable</code>	Enables preempt-mode for VRRPv3 for IPv4.

Table continues...

Variable	Value
<i>default ip vrrp <vrid> preempt-mode</i>	Sets the default preempt-mode value for VRRPv3 for IPv4.
<i>no ip vrrp <vrid> preempt-mode enable</i>	Disables preempt-mode for VRRPv3 for IPv4.

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.

If you have VRRP and IP routing protocols 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 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. 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.

Note:

The VRRP virtual IP address cannot be the same as the local IP address of the port or VLAN on which VRRP is enabled.

Important:

The switch, 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 the switch 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.

Note:

The VSP VRRP IP address responds only to ICMP-based traceroute requests. It does not respond to UDP-based traceroute requests.

Before you begin

- Assign an IP address to the interface.
- Enable VRRP globally.

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. <ul style="list-style-type: none"> • 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.

Modifying VRRP parameters for an interface

Before you begin

- You must enable VRRP on a brouter port or VLAN.

About this task

You can manage and configure VRRP parameters for the routing interface.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **VRRP**.
3. Click the **Interface** tab.

- 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.

- In the **Action** check box, select an option.
- 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: <ul style="list-style-type: none"> 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	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

Table continues...

Name	Description
	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: <ul style="list-style-type: none"> • 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	Configures the critical IP address for VRRP. This command specifies an IP interface on the local router, which is 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. * Note: In this context, <i>local</i> implies an address from the same VRF as the IP interface where VRRP is being configured.
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. 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.

Table continues...

Name	Description
FasterAdvIntervalEnable	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disable.

Configuring VRRP on a V3 interface

Perform this procedure to configure VRRP on a V3 interface on either a brouter port or a VLAN.

Before you begin

- Assign an IPv4 address to the interface
- Enable routing globally
- Do not configure RSMLT on the VLAN

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **VRRP**.
3. Click the **V3 Interface** tab.
4. Click **Insert**.
5. Beside the **IfIndex** field, click **Port** or **VLAN**.
6. Select a port or VLAN.
7. Click **OK**.
8. Type the virtual router ID.
9. Type the primary IP address.
10. Type the advertisement interval.
11. Click **Insert**.

Interface field descriptions

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

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router.

Table continues...

Name	Description
PrimaryIpAddr	Specifies the virtual address assigned to the VRRP.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Shows the state of the virtual router interface. The possible states are <ul style="list-style-type: none"> • 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 the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvInterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second.
MasterIpAddr	Specifies the IP address of the physical interface of the Master's virtual router.
UpTime	Indicates the time interval since this virtual router exited the INIT state.
CriticalIpAddr	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.
CriticalIpAddrEnabled	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
BackUpMaster	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the vIST. The default is disabled.
BackUpMasterState	Indicates if the Backup-Master is operational up. If the switch is in Master state but the Backup-Master is enabled, then the BACKUP MASTER STATE will be down.
FasterAdvIntervalEnabled	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disable.

Table continues...

Name	Description
FasterAdvInterval	Configures the interval between VRRP advertisement messages. The default is 200. Enter the values in multiples of 200 milliseconds.
PreemptMode	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master.
Action	Lists options to override the hold-down timer manually and force preemption: <ul style="list-style-type: none"> • none does not override the timer. • preemptHoldDownTimer preempts the timer. This parameter applies only if the holddown timer is active.
HoldDownTimer	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.
MasterAdvInterval	On the VRRPv3 master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

Configuring VRRPv3 Checksum

Perform this procedure to configure VRRPv3 checksum on either a brouter port or a VLAN.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **VRRP**.
3. Click the **V3 Checksum** tab.
4. Click **Insert**.
5. Beside the **IfIndex** field, click **Port** or **VLAN**.
6. Select a port or a VLAN.
7. Select a type of checksum computation.
8. Select a VRRP version.

- Click **Insert**.

V3 Checksum field descriptions

Use the data in the following table to use the **V3 Checksum** tab.

Name	Description
rclpConffIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
ChkSumComputation	Specifies the type of checksum computation, with Pseudo Header or without Pseudo Header.
VrrpVersion	Specifies the VRRP version; unspecified, version 2, or version 3.

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

- In the Device Physical View tab, select a port.
- In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
- Click **IP**.
- Click the **VRRP** tab.
- Click **Insert**.
- In the **Insert VRRP** dialog box, enable **FasterAdvIntervalEnable**.
- In the **FasterAdvInterval** field, enter a value. You must set this value using multiples of 200 milliseconds.
- Click **Insert**.

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 pane, expand the following folders: **Configuration > VLAN**.
2. Click **VLANS > Basic**.
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**.

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. With multicast virtualization, the Virtual Services Platform 4000 also functions as multiple virtual multicast routers. 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.

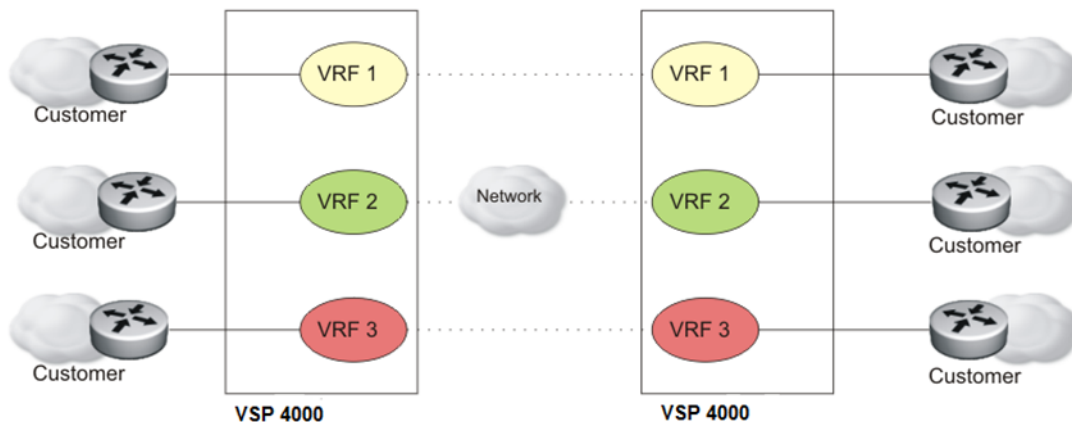


Figure 26: Multiple virtual routers in one system

One Avaya Virtual Services Platform 4000 Series 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 4000 Series 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:

- IP
- Internet Control Message Protocol (ICMP)
- Address Resolution Protocol (ARP)
- Static routes
- Default routes
- Route policies
- Virtual Router Redundancy Protocol (VRRP)
- Dynamic Host Configuration Protocol (DHCP), and BootStrap Protocol relay agent
- User Datagram Protocol (UDP) forwarding

Avaya Virtual Services Platform 4000 Series 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

VRF Lite interoperates with RFC 4364 and Layer 3 VPNs.

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 4000 Series.

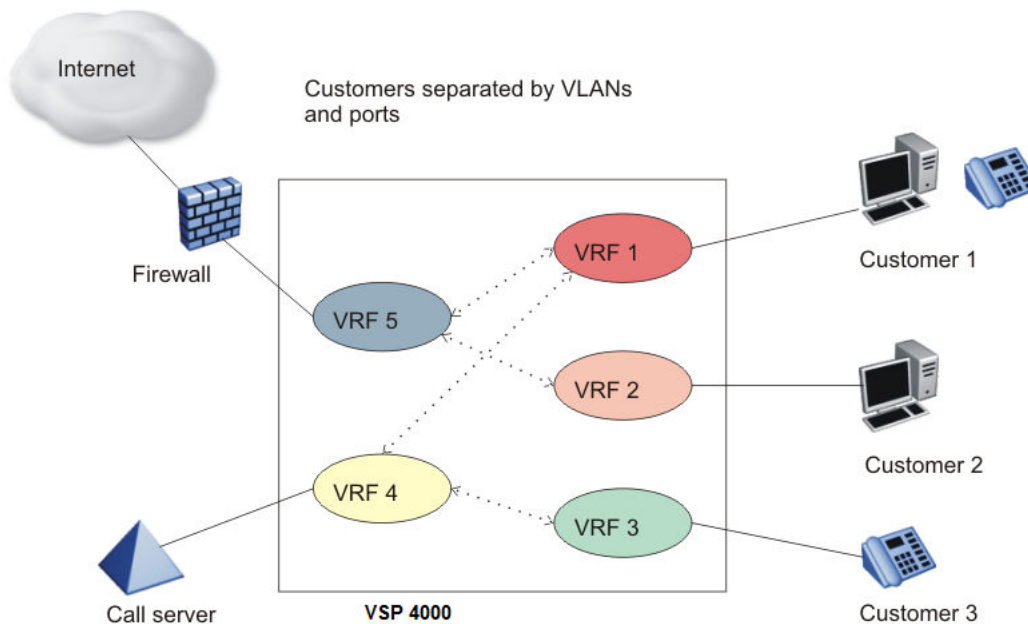


Figure 27: Inter-VRF forwarding

For more information about the latest VRF Lite scalability, see *Release Notes for VSP Operating System Software*, NN47227-401.

VRF Lite and inter-VRF route redistribution

The Avaya Virtual Services Platform 4000 Series supports three route redistribution functions:

- intra-VRF inter-protocol route redistribution (redistribution within the same VRF instance).
- inter-VRF inter-protocol redistribution (redistribution between two VRF instances).

- 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.

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. 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.

VRF Lite requirements

To use VRF Lite, you require the following hardware and software:

- Avaya Virtual Services Platform 4000 Series Software Release 3.0.0.0 or later

Port parameters 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 port 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 port parameters of the VRF. If a port belongs to more than one VRF instance, you cannot edit the port 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 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.

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.

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.

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.
- 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.
- Multiple VRFs on the same node can function in different autonomous systems.
- 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 56: Virtualized protocols and documentation

Virtualized protocol or feature	Where to find information
ARP	This document
Circuitless IP	This document
DHCP	This document
Route policies	This document
Route preferences	This document
Router Discovery	This document
Static routes	This document
User Datagram Protocol (UDP)	This document

Table continues...

Virtualized protocol or feature	Where to find information
VLAN	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 4000 Series, NN46251-500</i>
VRRP	This document

VRF Lite architecture examples

VRF Lite enables a router to act as many routers. This provides virtual traffic separation for each user and provides security. For example, you can use VRF Lite to:

- Provide different departments within a company with site-to-site connectivity as well as Internet access
- Provide centralized and shared access to data centers.

The following figure shows how VRF Lite can emulate VPNs.

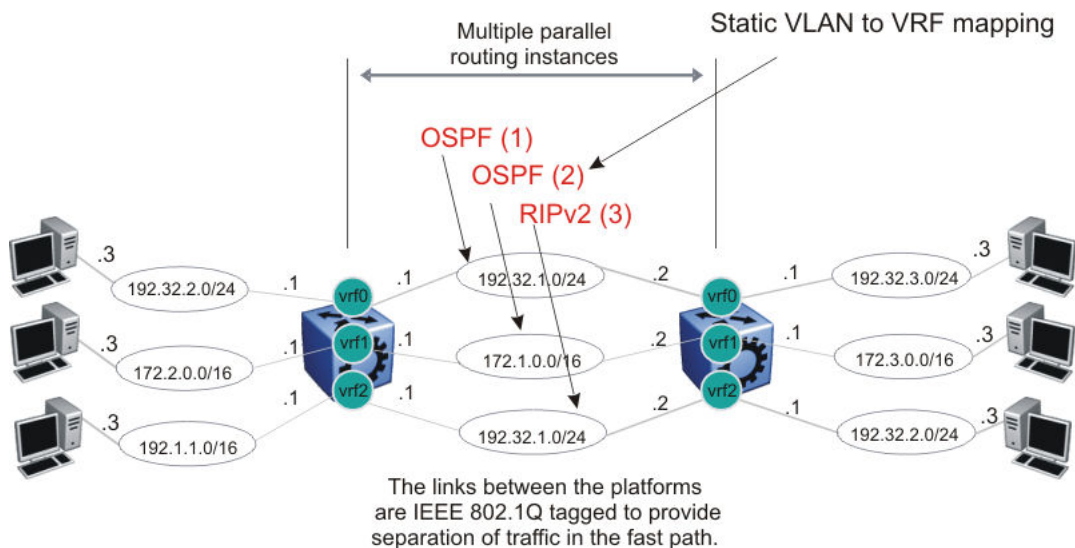


Figure 28: VRF Lite example

The following figure shows how VRFs can interconnect through an external firewall.

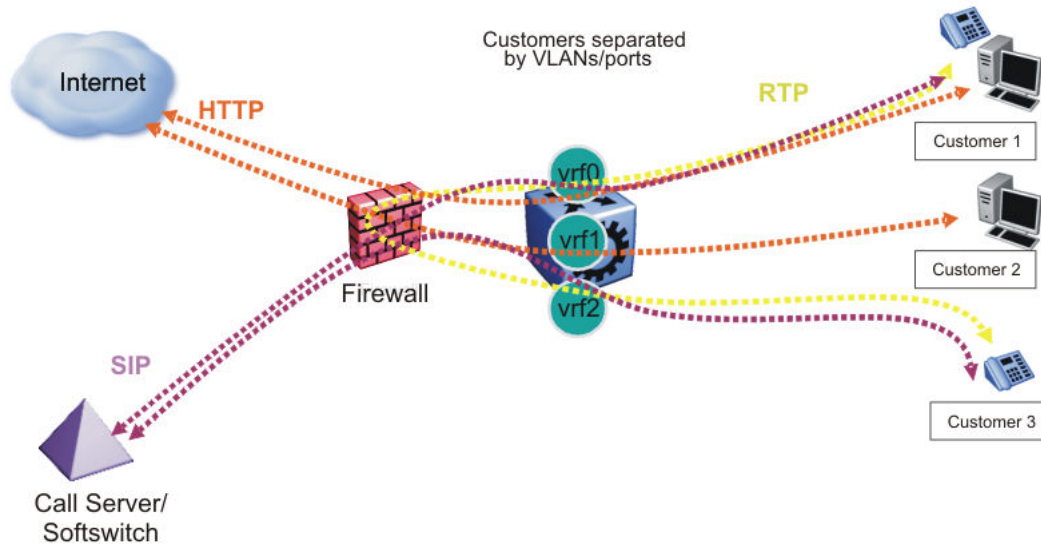


Figure 29: Inter-VRF forwarding based on external firewall

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, VoIP-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. Using the inter-VRF solution, you can use routing policies and static routes to inject IP subnets from one VRF instance to another, and filters to restrict access to certain protocols.

The following figure shows inter-VRF forwarding. In this solution, you can use routing policies to leak IP subnets from one VRF to another. You can use filters to restrict access to certain protocols. This configuration enables hub-and-spoke network designs, for example, for VoIP gateways.




Figure 30: Inter VRF communication, internal inter-VRF forwarding

Chapter 17: VRF Lite configuration using ACLI

Use Virtual Router and Forwarding (VRF) Lite to provide many virtual routers using one Virtual Services Platform 4000.

This section shows you how to configure a VRF instance and how to associate ports and VLANs with VRF instances.

 **Note:**

The prompt for the non-PowerPlus chassis is VSP-4850GTS. The prompt for the PowerPlus chassis is VSP-4850GTS-PWR+. The prompt for the Fiber box is VSP-4450 GSX. For consistency, this document uses the VSP-4850GTS prompt.

The following task flow shows you the sequence of procedures you perform to configure VRF Lite.




Figure 31: VRF Lite configuration procedures

*** Note:**

The prompt for the non-PowerPlus chassis is VSP-4850GTS. The prompt for the PowerPlus chassis is VSP-4850GTS-PWR+. The prompt for the Fiber box is VSP-4450 GSX. For consistency, this document uses the VSP-4850GTS prompt.

Creating a VRF instance

Before you begin

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

About this task

Create a VRF instance to provide a virtual routing interface for a user.

Procedure

1. Create a VRF instance and specify a VRF name:

```
ip vrf WORD<0-16>
```

2. Configure the maximum number of routes:

```
ip vrf WORD<0-16> max-routes <0-16000>
```

3. Enable max-routes traps:

```
ip vrf WORD<0-16> max-routes-trap enable
```

4. Enter Router Configuration mode:

```
router vrf WORD<0-16>
```

5. Ensure that the instance is configured correctly:

```
show ip vrf [WORD<0-16>]
```

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create a VRF instance and specify a VRF name:

```
VSP-4850GTS-PWR+:1(config)#ip vrf test1
```

Configure the maximum number of routes:

```
VSP-4850GTS-PWR+:1(config)#ip vrf test1 max-routes 12000
```

Enable max-routes traps:

```
VSP-4850GTS-PWR+:1(config)#router vrf test1 max-routes-trap enable
```

Enter Router Configuration mode:

```
VSP-4850GTS-PWR+:1(config)#router vrf test1
```

Ensure that the instance is configured correctly:

```
VSP-4850GTS-PWR+:1(router-vrf)#show ip vrf test1
```

Variable definitions

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

Table 57: Variable definitions

Variable	Value
max-routes <0-16000>	Specifies the maximum number of routes for the VRF. The default value is 10000, except for the Global Router, which is 16000.
max-routes-trap enable	Enables the sending of traps after the maximum number of routes is reached.
name <WORD 0-32>	Renames the VRF instance.
vrf-trap	Enables the device to send VRF-related traps.

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

Table 58: Variable definitions

Variable	Value
max-routes [vrfdids WORD<0-512>] [<WORD 0-16>]	Displays the maximum number of routes for the specified VRFs. <ul style="list-style-type: none"> vrfdids WORD<0-512> specifies a list of VRFs by VRF IDs. WORD<0-16> specifies a VRF by name.
vrfdids WORD<0-512>	Specifies a list of VRFs by VRF IDs.
WORD<0-16>	Specifies a VRF by name.

Example

Following is an output example for the `show ip vrf` command:

```
VSP-4850GTS-PWR+:1(config)#show ip vrf
```

```
=====
VRF INFORMATION
=====
VRF      OSPF      RIP      BGP      PIM      ARP
COUNT  COUNT    COUNT   COUNT   COUNT   COUNT
-----
1         0         0         0         0         0

VRF      VRF      VLAN  ARP
NAME    ID  OSPF  RIP   BGP   PIM   COUNT  COUNT
-----
GlobalRouter      0  FALSE FALSE FALSE FALSE  2      0

1 out of 1 Total Num of VRF Entries displayed.
```

Associating a VLAN or port with a VRF instance

Before you begin

- Ensure the VRF is already configured.
- You must log on to the Interface Configuration mode in ACLI.

About this task

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.

Procedure

1. Log on to VLAN Interface Configuration mode.
2. Associate the VLAN with a VRF instance:
3. Log on to GigabitEthernet Interface Configuration mode.
4. Associate a port with a VRF instance:

```
vrf WORD<0-16>
```

```
vrf <WORD 0-16>
```

Example

```
VSP-4850GTS-PWR+:1>enable
```

```
VSP-4850GTS-PWR+:1#configure terminal
```

Create a VRF named Two:

```
VSP-4850GTS-PWR+:1(config-if)#ip vrf Two
```

Create a VLAN of type byport:

```
VSP-4850GTS-PWR+:1(config-if)#vlan create 33 name vlan-30 type port-
mstp 0
```

Enter VLAN Interface Configuration mode:

```
VSP-4850GTS-PWR+:1(config-if)#interface vlan 33
```

Assign the VLAN to VRF Two:

```
VSP-4850GTS-PWR+:1(config-if)#vrf Two
```

Give the VLAN an IP address:

```
VSP-4850GTS-PWR+:1(config-if)#ip address 32.22.12.2 255.255.255.0
```

Enter VRF configuration mode:

```
VSP-4850GTS-PWR+:1(config-if)#router vrf Two
```

Variable definitions

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

Table 59: Variable definitions

Variable	Value
vrf WORD<0-16>	Specifies the VRF instance by name.

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 Fabric Connect on VSP Operating System Software*, NN47227-510.

Before you begin

- The VRF must exist.

Procedure

1. Specify the VRF to configure by logging on to VRF Router Configuration mode:

```
enable
configure terminal
router vrf WORD<0-16>
```

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<0-16>] [vrfids WORD<0-512>]
```

Example

From Global Configuration mode, log on to Router VRF Configuration mode:

```
VSP-4850GTS-PWR+:1(config)#router vrf red.
```

Create the IP VPN instance:

```
VSP-4850GTS-PWR+:1(router-vrf)#ipvpn
```

Enable IP VPN:

```
VSP-4850GTS-PWR+:1(router-vrf)#i-sid 100
```

Enable IP VPN:

```
VSP-4850GTS-PWR+:1(router-vrf)#ipvpn enable
```

```
VSP-4850GTS-PWR+:1#show ip ipvpn
VRF Name           : red
Ivpn-state         : enabled
I-sid              : 100
```

Variable definitions

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

Variable	Value
vrf <i>WORD</i> <0-16>	Specifies the VRF name.
vrfids <i>WORD</i> <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 <code>no</code> or <code>default</code> 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 4000 Series.

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.

Table continues...

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 15744.
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 use multiple RPs simultaneously. You can also 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. 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
DstVrflid	Specifies the destination VRF ID to use in the redistribution.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrflid	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 useful for OSPF and BGP only. The values are type1(1), and type2(2). The default is type2.
Subnets	Indicates that all the subnets must be advertised individually. The values are allow(1), and suppress(2). The default value is allow. This variable applies to OSPF only.

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 **BrouterVrflid** 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.
Vrflids	Identifies the set of VRF IDs to which this port belongs.
VrfCount	Shows the number of VRF instances associated with this port.
BrouterVrflid	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
ConfigNextAvailableVrflid	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.
RouterAddressType	Specifies the router address type of this VRF.
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.

Selecting and launching a VRF context view

About this task

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

Important:

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

Note:

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

Procedure

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 opens containing the selected VRF view

VRF field descriptions

Use the descriptions 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 port 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 Fabric Connect on VSP Operating System Software*, NN47227-510.

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
Vrflid	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.