



Troubleshooting Avaya Virtual Services Platform 9000

Release 4.0
NN46250-700
Issue 06.05
July 2015

© 2015 Avaya Inc.
All Rights Reserved.

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

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

Licenses

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

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The

applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

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 currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at <http://support.avaya.com/LicenseInfo> under the link "Heritage Nortel Products" or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants You a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

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

Third Party Components

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

Preventing Toll Fraud

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

result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

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

Downloading Documentation

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

Contact Avaya Support

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

Contents

Chapter 1: Introduction	10
Purpose.....	10
Related resources.....	10
Documentation.....	10
Training.....	10
Viewing Avaya Mentor videos.....	10
Support.....	11
Searching a documentation collection.....	11
Chapter 2: New in this release	13
Features.....	13
Other changes.....	17
Chapter 3: Data collection required for Technical Support cases	18
Data collection for an outage.....	18
Collecting data before you restart.....	18
Displaying current patch information.....	24
Collecting data after you restart.....	25
Data collection for non outage problems.....	26
Gathering critical information.....	26
Collecting data.....	26
Chapter 4: Troubleshooting fundamentals	31
Troubleshooting planning fundamentals.....	31
Proper installation and routine maintenance.....	31
Network configuration.....	32
Normal behavior on the network.....	33
Troubleshooting fundamentals.....	34
Connectivity problems.....	34
Routing table problems.....	35
LED indications of problems.....	35
Cable connection problems.....	37
QSFP+ authentication error.....	38
Alarm database.....	38
Troubleshooting tool fundamentals.....	39
Troubleshooting overview.....	39
Digital Diagnostic Monitoring.....	41
Port mirroring.....	42
Remote mirroring.....	45
Packet Capture Tool.....	49
Flight Recorder.....	54
General diagnostic tools.....	55

Chapter 5: Log and trap information	57
Log and trap fundamentals.....	57
Simple Network Management Protocol.....	57
Overview of traps and logs.....	58
Log message format.....	59
Log files.....	61
Log file transfer.....	62
Log configuration using ACLI.....	63
Configuring a UNIX system log and syslog host.....	64
Configuring logging.....	67
Configuring the remote host address for log transfer.....	68
Configuring system logging to external storage.....	69
Configuring system message control.....	70
Extending system message control.....	71
Viewing logs.....	72
Configuring ACLI logging.....	75
Log configuration using EDM.....	77
Configuring the system log.....	77
Configuring the system log table.....	78
SNMP trap configuration using ACLI.....	79
Configuring an SNMP host.....	80
Configuring an SNMP notify filter table.....	81
Configuring SNMP interfaces.....	82
Enabling SNMP trap logging.....	85
SNMP trap configuration using EDM.....	87
Configuring an SNMP host target address.....	87
Configuring target table parameters.....	89
Configuring an SNMP notify table.....	90
Configuring SNMP notify filter profiles.....	91
Configuring SNMP notify filter profile table parameters.....	91
Enabling SNMP trap logging.....	92
Viewing the trap sender table.....	93
Chapter 6: Hardware troubleshooting	94
Troubleshooting module failure.....	94
Troubleshooting CP start failure.....	95
Removing external storage devices from the CP module.....	95
Troubleshooting USB viewing problems.....	97
Troubleshooting USB writing problems.....	99
Troubleshooting external Compact Flash viewing problems.....	99
Using trace to diagnose hardware problems.....	100
Chapter 7: Software troubleshooting	102
Software troubleshooting.....	102
Failure to read the configuration file.....	102

No web management interface access to a device.....	102
Cannot enable encryption.....	103
Debug files.....	103
Software download.....	104
Downloading the software.....	104
Downloading Avaya Virtual Services Platform 9000 documentation.....	105
Software troubleshooting tool configuration using the ACLI.....	105
Using ACLI for troubleshooting.....	105
Using software record dumps.....	109
Using trace to diagnose problems.....	110
Using trace to diagnose Ipv6 problems.....	113
Viewing and clearing the alarm database.....	115
Viewing and deleting debug files.....	116
Configuring port mirroring.....	121
Configuring global mirroring actions with an ACL.....	125
Configuring ACE actions to mirror.....	126
Configuring Layer 2 remote mirroring.....	127
Accessing the secondary CPU.....	130
Configuring PCAP global parameters.....	130
Enabling PCAP on a port.....	133
Configuring PCAP capture filters.....	134
Using the captured packet dump.....	138
Copying captured packets to a remote machine.....	138
Resetting the PCAP DRAM buffer.....	139
Clearing ARP information for an interface.....	140
Flushing routing, MAC, and ARP tables for an interface.....	141
Pinging an IP device.....	142
Running a traceroute test.....	143
Showing SNMP logs.....	144
Using trace to examine IS-IS control packets.....	145
Software troubleshooting tool configuration using EDM.....	151
Flushing routing tables by VLAN.....	151
Flushing routing tables by port.....	151
Configuring port mirroring.....	152
Configuring Layer 2 remote mirroring.....	153
Configuring ACLs for mirroring.....	154
Configuring ACEs for mirroring.....	156
Configuring PCAP globally.....	159
Configuring PCAP on a port.....	160
Configuring PCAP filters.....	161
Configuring advanced PCAP filters.....	163
Running a ping test.....	165
Viewing ping results.....	167

Viewing ping probe history.....	168
Running a traceroute test.....	168
Viewing traceroute results.....	170
Viewing the traceroute history.....	171
Chapter 8: Layer 1 troubleshooting.....	173
Troubleshooting fiber optic links.....	173
Resetting a QSFP+ transceiver.....	174
Chapter 9: Layer 2 troubleshooting.....	176
Troubleshooting IST failure.....	176
Troubleshooting BPDU Filtering.....	180
No packets received on the port.....	180
SNMP trap not received.....	182
Chapter 10: Connectivity Fault Management.....	184
CFM fundamentals.....	184
Autogenerated CFM and explicitly configured CFM.....	185
Maintenance Domain (MD).....	187
Maintenance Association (MA).....	188
Maintenance endpoints (MEP).....	189
Maintenance domain intermediate points (MIP).....	189
Fault verification.....	189
LBM message.....	190
I2 ping.....	190
Fault isolation.....	191
Link trace message.....	191
I2 traceroute.....	192
I2 tracetree.....	194
Layer 2 tracemroute.....	194
Nodal MPs.....	194
Configuration considerations.....	195
CFM configuration using ACLI.....	196
Autogenerated CFM.....	196
Configuring explicit CFM.....	201
Triggering a loopback test (LBM).....	207
Triggering linktrace (LTM).....	208
Triggering a Layer 2 ping.....	210
Triggering a Layer 2 traceroute.....	212
Triggering a Layer 2 tracetree.....	214
Triggering a Layer 2 tracemroute.....	216
Using trace CFM to diagnose problems.....	218
Using trace SPBM to diagnose problems.....	221
CFM configuration using EDM.....	224
Autogenerated CFM.....	224
Configuring explicit CFM.....	227

Configuring Layer 2 ping.....	232
Initiating a Layer 2 traceroute.....	234
Viewing Layer 2 traceroute results.....	237
Configuring Layer 2 IP ping.....	238
Viewing Layer 2 IP ping results.....	241
Configuring Layer 2 IP traceroute.....	242
Viewing Layer 2 IP traceroute results.....	244
Triggering a loopback test.....	246
Triggering linktrace.....	248
Viewing linktrace results.....	250
Configuring Layer 2 tracetable.....	252
Viewing Layer 2 tracetable results.....	255
Configuring Layer 2 trace multicast route on a VLAN.....	256
Configuring Layer 2 tracemroute on a VRF.....	258
Viewing Layer 2 trace multicast route results.....	260
Chapter 11: Upper layer troubleshooting.....	262
Troubleshooting SNMP.....	262
Troubleshooting DHCP.....	263
Troubleshooting DHCP Relay.....	263
Troubleshooting client connection to the DHCP server.....	266
Troubleshooting IPv6 DHCP Relay.....	267
IPv6 DHCP Relay switch side troubleshooting.....	267
IPv6 DHCP Relay server side troubleshooting.....	268
IPv6 DHCP Relay client side troubleshooting.....	268
Enabling trace messages for IPv6 DHCP Relay.....	269
Troubleshooting IPv6 VRRP.....	269
VRRP transitions.....	269
Enabling trace messages for IPv6 VRRP troubleshooting.....	270
Risks associated with enabling trace messages.....	272
VRRP with higher priority running as backup.....	272
Troubleshooting RSMLT.....	273
RSMLT configuration considerations.....	273
RSMLT peers not up.....	273
Enabling trace messages for RSMLT troubleshooting.....	274
Troubleshooting IPv6 connectivity loss.....	275
Troubleshooting TACACS+.....	275
Unable to log on using Telnet or rlogin.....	276
Unable to log on using SSH.....	280
Unable to log on by any means (Telnet, rlogin, or SSH).....	281
Administrator unable to obtain accounting information from the TACACS+ server.....	285
Trap server cannot receive trap packets from the VSP device.....	285
Troubleshooting TACACS+ problems.....	287
Troubleshooting client registration.....	288

Chapter 12: Unicast routing troubleshooting	289
Using BGP debugging commands.....	289
Troubleshooting licensed routing protocols.....	291
Viewing OSPF errors.....	292
Viewing OSPF neighbor state problems.....	294
Troubleshooting OSPF Init state problems.....	295
Troubleshooting OSPF ExStart/Exchange problems.....	296
Chapter 13: Multicast troubleshooting	298
Multicast feature troubleshooting.....	298
Troubleshooting IGMP Layer 2 Querier.....	298
Troubleshooting IGMPv3 backwards compatibility.....	302
Multicast routing troubleshooting using ACLI.....	304
Viewing IGMP interface information.....	304
Viewing multicast group trace information for IGMP snoop.....	307
Viewing IGMP group information.....	309
Showing the hardware resource usage.....	311
Using PIM debugging commands.....	312
Determining the protocol configured on the added VLAN.....	314
Determining the data stream learned with IP multicast over SPBM on the VLAN.....	317
Displaying the SPBM multicast database.....	319
Troubleshooting IP multicast over SPBM for Layer 2 VSNs.....	320
Troubleshooting IP multicast over SPBM for Layer 3 VSNs.....	322
Troubleshooting IP multicast over SPBM for IP Shortcuts.....	324
Defining the IS-IS trace flag for IP multicast.....	325
Multicast routing troubleshooting using EDM.....	325
Viewing IGMP interface information.....	325
Viewing IGMP snoop trace information.....	328
Viewing IGMP group information.....	329
Viewing multicast group sources.....	330
Viewing multicast routes by egress VLAN.....	330
Determining the data stream learned when IP multicast over SPBM is configured on the VLAN.....	331
Showing the SPBM multicast database.....	332
Chapter 14: Safety Messages	333
Glossary	337

Chapter 1: Introduction

Purpose

Troubleshooting describes common problems and error messages, provides information about traps and command logging, and provides techniques you can use to resolve common problems.

Troubleshooting also provides information about troubleshooting tools: for example, port and remote mirroring.

Related resources

Documentation

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

Training

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

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

Viewing Avaya Mentor videos

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

About this task

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

Procedure

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

 **Note:**

Videos are not available for all products.

Support

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

Searching a documentation collection

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

Before you begin

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

Procedure

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

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

Chapter 2: New in this release

The following sections detail what is new in *Troubleshooting Avaya Virtual Services Platform 9000*, NN46250-700, for Release 4.0.

Features

See the following sections for information about feature changes.

Digital Diagnostic Interface (DDI)

Release 4.0 updates information on Digital Diagnostic Monitoring. Digital Diagnostic Interface (DDI) module information for RxPower can output false alerts. The MSA (Multi-source Agreement) between manufacturers of SFP and SFP+ devices specifies a +/- 3dB accuracy tolerance for optical power measurements.

To minimize false warnings or alarms due to this inaccuracy, the thresholds for low and high TxPower and for low RxPower are offset by this tolerance. High RxPower thresholds are not offset due to the potential for receiver saturation and damage that can result over the long-term, however, this increases the possibility of false alerts.

If high RxPower alerts occur, but the link operates normally, consider this tolerance. If the link fails to operate, consider the possibility that the optical receiver is being over-driven, and attempt to correct the condition.

For more information, see [Digital Diagnostic Monitoring](#) on page 41.

Digital Diagnostic Monitoring (DDM)

Release 4.0 updates the Digital Diagnostic Monitoring (DDM) feature in the following ways:

- Digital Diagnostic Monitoring (DDM) is now enabled by default. The `pluggable-optical-module ddm-monitor` and `no pluggable-optical-module ddm-monitor` commands to enable and disable DDM are no longer available.
- The `pluggable-optical-module ddm-alarm-portdown` command is also no longer available.
- Release 4.0 also updates the `pluggable-optical-module ddm-monitor-interval <10..40>`. The DDM monitor interval is now 10–40 seconds. The default is 10 seconds.
- The device sends traps if you enable the DDM traps send feature using the `pluggable-optical-module ddm-traps-send` command. The device always generates logs no matter if the DDM traps send feature is enabled or disabled.

- For the **show pluggable-optical-modules basic** command, the device reports qualified optics as Avaya in the type field. The device reports non-qualified best-effort optics as a different manufacturer in the type field. Unsupported optics display as unsupported in the type field, and do not operate in the system.

Virtual Services Platform 9000 continues to support older configurations prior to Release 4.0.

For more information, see:

- [Digital Diagnostic Monitoring](#) on page 41.

Layer 2 ping

Syntax corrected in the procedure for the **12 ping** command. For more information, see [Triggering a Layer 2 ping](#) on page 210.

Licensed feature content

Release 4.0 adds support for one million IPv4 FIB routes under the premier license, if you use second generation modules, and configure the Virtual Services Platform 9000 to use second generation modules.

For more information on licensing, see: [Job aid](#) on page 292.

Log and trap information

Release 4.0 restructures the log and trap information section of the document. Release 4.0 adds an additional ACLI procedure named Viewing the trap sender table. For more information, see:

- [Viewing the trap sender table](#) on page 93.

Release 4.0 renames the EDM procedure originally called Enabling authentication traps to Enabling SNMP trap logging to match the ACLI procedure name. For more information, see:

- [Enabling SNMP trap logging](#) on page 92.

Release 4.0 updates parameters in the **snmp-server** command variable definition table. For more information, see:

- [Variable definitions](#) on page 83.

Maintenance Domain name and Maintenance Association name

Release 4.0 updates the variable range to <1–22> for Maintenance Domain (MD) name and Maintenance Association (MA) name for Connectivity Fault Management (CFM). For more information, see:

- [Configuring CFM MD](#) on page 201.
- [Configuring CFM MA](#) on page 202.
- [Configuring CFM MEP](#) on page 205.
- [Triggering a loopback test \(LBM\)](#) on page 207.
- [Triggering linktrace \(LTM\)](#) on page 208.

Minor update to syntax

Release 4.0 made a minor correction to syntax in the following procedure:

- [Configuring port mirroring](#) on page 121.

PCAP dump error

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

For more information, see:

- [Configuring PCAP global parameters](#) on page 130.
- [Enabling PCAP on a port](#) on page 133.
- [Configuring PCAP capture filters](#) on page 134.
- [Using the captured packet dump](#) on page 138.
- [Resetting the PCAP DRAM buffer](#) on page 139.
- [Configuring PCAP globally](#) on page 159.
- [Configuring PCAP on a port](#) on page 160.
- [Configuring PCAP filters](#) on page 161.
- [Configuring advanced PCAP filters](#) on page 163.

PCAP capture filters

Release 4.0 updates the `user-defined<0–9600> WORD<1-50>` parameter in the `pcap capture filter <1–1000> user-defined <0–9600> WORD<1–50>` command to the variable `WORD<1–50>`. For more information, see [Variable definitions](#) on page 135.

QSFP+ support

Release 4.0.1 adds support for quad small form factor pluggable plus (QSFP+) for use with the 9012QQ-2 module. For more information, see:

- [Cable connection problems](#) on page 37.
- [QSFP+ authentication error](#) on page 38.
- [Digital Diagnostic Monitoring](#) on page 41.
- [Troubleshooting fiber optic links](#) on page 173.
- [Resetting a QSFP+ transceiver](#) on page 174.

Removed trace auto command

In Release 4.0, the content related to the `trace auto` command was removed because this command is not supported on Virtual Services Platform 9000. This command applies only to ERS 8800.

Router node name

The Virtual Services Platform 9000 does not support the `routernodename` option for C-VLANs with the `12 ping` and `12 traceroute` commands. For more information, see:

- [Variable definitions](#) on page 211.
- [Variable definitions](#) on page 214.

SNMP Params field

Release 4.0 updates the SNMP **Params** field. For more information, see [Configuring an SNMP host target address](#) on page 87.

Statistics procedures moved

Release 4.0 removes ACLI and EDM procedures to enable multicast routing process statistics. The information now exists in the *Monitoring Performance on Avaya Virtual Services Platform 9000*, NN46250-701, document with other statistics information.

Unprintable characters in log and trace messages

Unprintable characters in log and trace messages are encoded in C language hexadecimal notation in the string and surrounded with the `<UNPRINTABLE>...<UNPRINTABLE>` tag. For example, the device converts a 4-byte sequence of 0x01, 0x02, 0xfe, 0xff to `\x01\x02\xfe\xff`, and adds the `<UNPRINTABLE>...<UNPRINTABLE>` tag to the log or trace message. The following example displays how the message appears:

```
CP1 [07/31/14 15:53:36.225] 0x000e4609 00000000 GlobalRouter SW INFO <UNPRINTABLE>
rlogind: session 0 *IN USE* via user: \xc3\xb4\xb6\x8d\x0e\xb9+\xb1\xa2aLO~o\xfb3m\xf9\x8c
\x05g\xd2.\xc9pSWP\xc5\xd9 ip 10.139.82.29 <UNPRINTABLE>
```

For more information, see

- [Log message format](#) on page 59.
- [Viewing logs](#) on page 72.
- [Configuring ACLI logging](#) on page 75.

Update to show ip route command output

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

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

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

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

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

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

- [Collecting data before you restart](#) on page 18.
- [Unable to log on by any means \(Telnet, rlogin, or SSH\)](#) on page 281.

Update to show logging file CPU parameter

Release 4.0 updates the `show logging file CPU` parameter variable to `WORD<0-100>`. For more information, see [Viewing logs](#) on page 72.

Update to the show tech command

Release 4.0 updates the output for the `show tech` command. The update adds Field Programmable Gate Array (FPGA) temperature sensor and Fabric Adapter (FA) temperature sensor readings to the output for second generation modules. For more information, see:

- [Normal behavior on the network](#) on page 33.
- [Using ACLI for troubleshooting](#) on page 105.

Update to show mlt command output

Release 4.0 updates the output for the `show mlt <1-512>` command. For more information, see [Troubleshooting IST failure](#) on page 176.

Update to show running-config

Release 4.0 sorts the `show running-config module` command parameters in alphabetical order. For more information, see [Using ACLI for troubleshooting](#) on page 105.

Update to SNMP parameter information

Minor updates were made to SNMP parameter information. For more information, see:

- [Variable definitions](#) on page 81.
- [Variable definitions](#) on page 83.
- [Variable definitions](#) on page 182.

Other changes

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

Document title

Release 4.0 updates the document title to *Troubleshooting Avaya Virtual Services Platform 9000, NN46250–700*, from *Avaya Virtual Services Platform Troubleshooting, NN46250–700*.

Downloading the software

Release 4.0 changes the location of software downloads. For more information, see [Downloading the software](#) on page 104.

Traps reference

Release 4.0 moves traps reference information into its own document. For more information, see *Traps Reference for Avaya Virtual Services Platform 9000, NN46250-704*.

Chapter 3: Data collection required for Technical Support cases

Use the following sections to learn about how to gather information before you contact Avaya for technical support.

Data collection for an outage

Perform the following data collection procedures when Avaya Virtual Services Platform 9000 is in an outage condition and you require Avaya Technical Support to perform a root cause analysis.

Collecting data before you restart

Perform this procedure before you restart the chassis, or individual modules (Control Processor (CP) or interface).

About this task

Flight recorder

Flight recorder stores history and the current state for various kernel, system, and application data with minimum overhead to execute. After a debug crash, you can access this data to help you troubleshoot. You can take the PMEM snapshot, always-on trace, and archive snapshot with the `flight-recorder all {slot[-slot] [, ...]}` command.

Chassis information

Use the `show fulltech` commands to capture the current state of the chassis. The `show fulltech [file WORD<1-99>]` command runs all show commands and allows you to capture the output to a file.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Capture Flight Recorder trace information for each interface module that has active ports in the network, and for the Master and Backup CP modules:

```
flight-recorder all {slot[-slot][,...]}
```

This command executes three separate commands:

- `flight-recorder snapshot`
- `flight-recorder trace`
- `flight-recorder archive`

*** Note:**

Avaya recommends you run the `flight-recorder all` command first to capture the state of the system. If you run the `show fulltech` or other commands before the `flight-recorder all` command, the capture displays the state after you execute the `show fulltech` command.

The generated .tar file includes the following types of files:

Name	Type	Modified	Size	Ratio	Packed	Path
version.cfg	CFG File	2/18/2011 8:14 AM	146	0%	146	
trace.20110218121419.1.txt	Text Document	2/18/2011 8:14 AM	33,010,...	0%	33,01...	
pmem.20110218121414.1.bin.gz	WinZip File	2/18/2011 8:14 AM	389,902	0%	389,902	
messages	File	2/18/2011 8:14 AM	38,096	0%	38,096	
log.a1700001.226.gz	WinZip File	2/18/2011 8:14 AM	97,222	0%	97,222	
config.cfg	CFG File	2/18/2011 8:14 AM	17,378	0%	17,378	
catcs.txt	Text Document	2/18/2011 8:14 AM	1,669	0%	1,669	
archive.sh	SH File	2/18/2011 8:14 AM	635	0%	635	

3. Capture the current state of the chassis:

```
terminal more disable
show fulltech [file WORD<1-99>]
```

*** Note:**

Issue the `show fulltech [file WORD<1-99>]` command a second time, a few seconds apart from the first execution, which helps give a picture of what counters increment.

4. Repeat step 3 on page 19 on the IST peer VSP node, and if time permits, on other neighbor nodes to the VSP 9000 that exhibits the problem.
5. Display all files under `/intflash` and `/extflash` on both CPs:

```
terminal more disable
ls -r
cd /extflash
ls -r
cd /mnt/intflash
ls -r
```

Data collection required for Technical Support cases

```
cd /mnt/extflash
ls -r
terminal more enable
```

6. If you suspect issues with routing, also collect information on the VRF specific ARP table:

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

7. Display the route table for each VRF:

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

8. Reset the chassis:

```
reset -y
```

9. Continue with [Collecting data after you restart](#) on page 25.

Example

The following example shows output of the `flight-recorder all` command for slot 1 only. You must use this command for all active slots as identified in the procedure steps.

```
VSP-9012:1#flight-recorder all 1
Processing Flight-recorder snapshot for 1 ....

Flight-recorder snapshot for slot 1 complete, filename is /intflash/PMEM/1/pmem.
20111019114431.1.bin.gz.

Processing Flight-recorder trace for 1 ....

Flight-recorder trace for slot 1 complete, filename is /intflash/flrec/1/trace.2
0111019114434.1.txt.

Processing Flight-recorder archive for slot 1 ....

Flight-recorder archive for slot 1 complete, filename is /intflash/archive/1/arch
hive.20111019114446.1.tar.
```

The following example shows output of the `flight-recorder all` command for all module types and all slots in the chassis:

```
VSP-9012:1>enable
VSP-9012:#flight-recorder all all
Processing Flight-recorder snapshot for 1 ....

Flight-recorder snapshot for slot 1 complete, filename is /intflash/PMEM/1/pmem.
20111019113929.1.bin.gz.

Processing Flight-recorder trace for 1 ....

Flight-recorder trace for slot 1 complete, filename is /intflash/flrec/1/trace.2
0111019113931.1.txt.

Processing Flight-recorder archive for slot 1 ....

Flight-recorder archive for slot 1 complete, filename is /intflash/archive/1/arch
hive.20111019113944.1.tar.

Processing Flight-recorder snapshot for 4 ....

Flight-recorder snapshot for slot 4 complete, filename is /intflash/PMEM/4/pmem.
```

```

20111019113948.4.bin.gz.
Processing Flight-recorder trace for 4 ....
Flight-recorder trace for slot 4 complete, filename is /intflash/flrec/4/trace.20111019113952.4.txt.
Processing Flight-recorder archive for slot 4 ....
Flight-recorder archive for slot 4 complete, filename is /intflash/archive/4/archive.20111019113956.4.tar.
Processing Flight-recorder snapshot for 11 ....
Flight-recorder snapshot for slot 11 complete, filename is /intflash/PMEM/11/pmem.20111019114006.11.bin.gz.
Processing Flight-recorder trace for 11 ....
Flight-recorder trace for slot 11 complete, filename is /intflash/flrec/11/trace.20111019114010.11.txt.
Processing Flight-recorder archive for slot 11 ....
Flight-recorder archive for slot 11 complete, filename is /intflash/archive/11/archive.20111019114014.11.tar.
Processing Flight-recorder snapshot for SF4 ....
Flight-recorder snapshot for slot SF4 complete, filename is /intflash/PMEM/SF4/pmem.20111019114030.SF4.bin.gz.
Processing Flight-recorder trace for SF4 ....
Flight-recorder trace for slot SF4 complete, filename is /intflash/flrec/SF4/trace.20111019114038.SF4.txt.
Processing Flight-recorder archive for slot SF4 ....
Flight-recorder archive for slot SF4 complete, filename is /intflash/archive/SF4/archive.20111019114042.SF4.tar.

```

The following example displays the **show fulltech** command saved to a certain file:

```

VSP-9012:1#show fulltech file olivertech.txt
Saving fulltech results into file..
Dump fulltech result to file olivertech.txt successfully.

```

Display all show commands:

```

VSP-9012:1(config)#show fulltech
*****
          Show Fulltech
          Time           : Thu Jul 11 20:29:33 2013 EDT
*****

```

Data collection required for Technical Support cases

```
Command:[1] [ show access-policy by-mac ]
-----
      default-action      :    allow
      MAC Address         Action
      =====

```

```
Command:[2] [ show access-policy snmp-group ]
-----
      snmpv3-groups :
Policy 1 snmpv3-groups:
      Group Name          Snmp-Model
Policy 2 snmpv3-groups:
      Group Name          Snmp-Model

```

```
Command:[3] [ show access-policy ]
-----
AccessPolicyEnable: off
      Id: 1
      Name: default
PolicyEnable: true
      Mode: allow
      Service: ftp|http|telnet|ssh
      Precedence: 128
      NetAddrType: any
      NetAddr: N/A
      NetMask: N/A
TrustedHostAddr: N/A
TrustedHostUserName: none
      AccessLevel: readOnly
      AccessStrict: false
      Usage: 0
      Id: 2
      Name: snmpv3
PolicyEnable: true
      Mode: allow
      Service: snmpv3
      Precedence: 10
      NetAddrType: any
      NetAddr: N/A
      NetMask: N/A
TrustedHostAddr: N/A
TrustedHostUserName:
      AccessLevel: readOnly
      AccessStrict: false
      Usage: 0
--More-- (q = quit)

```

Display files in a directory:

```
VSP-9012:1(config)#ls -r
drwxr-xr-x 19 0 0 4096 Jul 2 08:22 ./
drwxrwxr-x 21 0 0 0 Jul 2 08:22 ../
drwx----- 2 0 0 16384 Feb 28 2010 lost+found/
drwxr-xr-x 6 0 0 4096 Jul 2 08:21 release/
drwx---r-x 2 0 0 4096 Feb 28 2010 common/
drwxr-xr-x 20 0 0 4096 Feb 28 2010 PMEM/
drwxr-xr-x 20 0 0 4096 May 19 2010 coreFiles/
drwxr-xr-x 2 0 0 4096 Jul 2 08:22 bootlog/
drwxr-xr-x 14 0 0 4096 Feb 28 2010 дума/
drwxr-xr-x 2 0 0 4096 Jul 2 08:23 sysInfo/
-rwx----- 1 0 0 6007 Mar 18 2010 jnk.jnk
-rw----- 1 0 0 117 Jul 2 08:21 .00000-11111.enc
--w----- 1 0 0 8722 May 13 2010 alarmLog.080
drwxr-xr-x 3 0 0 4096 Oct 7 2011 trace/
-rw-r--r-- 1 0 0 11 Jul 2 08:22 engboot
-rw----- 1 0 0 317 Jul 2 08:22 .shadov.txt
--w----- 1 0 0 15976765 May 13 2010 ale00001.000
-rw----- 1 0 0 8 Jul 2 11:08 .ospf_md5key.txt
-rw----- 1 0 0 0 Jul 2 11:08 ospf_vrfif_md5key.txt
-rw----- 1 0 0 8 Jul 2 11:08 ospf_vrfvif_md5key.txt
--w----- 1 0 0 3169 Jul 2 11:08 snmp_comm.txt
--w----- 1 0 0 12 Jul 2 11:08 snmp_usm.txt
--w----- 1 0 0 4385 May 13 2010 alarmLog.081
```

Display ARP information for a particular VRF:

```
VSP-9012:1#show ip arp vrf wan
```

```
=====
IP Arp - VRF wan
=====
IP_ADDRESS          MAC_ADDRESS          VLAN    PORT  TYPE    TTL(10 Sec)  TUNNEL
-----
192.0.2.149         00:00:0c:07:ac:c1    3901    MLT 512  DYNAMIC 1828
192.0.2.148         00:16:c7:af:e4:d8    3901    MLT 512  DYNAMIC 1254
192.0.2.151         ff:ff:ff:ff:ff:ff    3901    -      LOCAL   2160
192.0.2.145         00:1b:4f:60:aa:02    3901    -      LOCAL   2160
192.0.2.147         00:16:9d:4f:2a:10    3901    3/4    DYNAMIC 1236
192.0.2.146         2c:f4:c5:95:0a:02    3901    MLT 512  DYNAMIC 1958
192.0.2.153         00:1b:4f:60:aa:04    3903    -      LOCAL   2160
192.0.2.155         ff:ff:ff:ff:ff:ff    3903    -      LOCAL   2160
192.0.2.154         2c:f4:c5:95:0a:04    3903    MLT 512  DYNAMIC 2159
192.0.2.161         00:1b:4f:60:aa:05    3905    -      LOCAL   2160
192.0.2.163         ff:ff:ff:ff:ff:ff    3905    -      LOCAL   2160
192.0.2.162         00:16:9d:4f:2a:10    3905    3/4    DYNAMIC 2160
=====
IP Arp Extn - VRF wan
=====
MULTICAST-MAC-FLOODING  AGING (Minutes)      ARP-THRESHOLD
-----
N/A                      360                  N/A
```

```
12 out of 126 ARP entries displayed
```

Display the route table for each VRF:

```
VSP-9012:1#show ip route vrf test
```

```
=====
test                                     IP Route - VRF
=====
```

DST	MASK	NEXT	NH VRF/ISID	INTER COST FACE	PROT	AGE	TYPE	PRF
192.0.2.1	255.255.255.255	1.1.1.5	GlobalRouter	0 0	ISIS	0	IB	200
198.51.100.2	255.255.255.0	5.7.1.1	-	1 7	LOC	0	DB	0
192.0.3.16	255.255.255.0	VSP13	GlobalRouter	10 1000	ISIS	0	IBSV	7

Displaying current patch information

Use this procedure to display and gather current patch information.

Avaya requests this information when reproducing a field area that you have reported. You can also use the procedure to confirm a patch application loaded properly.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display current patch information:
show software patch all

*** Note:**

You must scroll to the end to see the patch information.

Example

```
VSP-9012:1>show software patch all
```

```
Patch Info:
-----
=====
Patch Status Information (All)
=====
```

Identifier	Type	Software	Status	Title
R00905765A	rst	3.1.0.0.GA	av	VSP9000: Pings to NLB Server Virtual IP
Fail after 15 Minutes T00896935A	hls	3.1.0.0.GA	av	VSP9000: FTP Authentication Fails with RADIUS
T00907580A	hls	3.1.0.0.GA	av	VSP9000:NLB: Client PC cannot connect to


```

NLB Services
T00942090A          hls  3.1.0.0.GA  av  VSP9000:SSH: CP Crash Due to SSH

Type:  rst = reset, hls = hitless, htfl = hitful
Status: ap = applied, ca = candidate, av = available, un = unknown

```

Collecting data after you restart

Perform this procedure after you restart the affected chassis, CP module, or interface module.

Procedure

1. Use FTP to transfer the following information:

- configuration files from each chassis — Configuration files are stored on the internal flash at `/intflash/`.
- log files from each chassis — Log files are stored on the external flash at `/extflash/`. If external flash does not exist, the system raises an alarm, and then logs are stored to internal flash instead. The log file is named using the format `log.xxxxxxxx.sss` and the alarm log is named `alarmLog`.
- generated archive files for each slot — The archive files are stored on the external flash at `/extflash/archive/[slot#]`. If external flash does not exist, the files are stored on the internal flash at `/intflash/archive/[slot]`. See [Collecting data before you restart](#) on page 18 for example output that shows how to identify the location and filename of the archive files.

*** Note:**

Some exceptions can result in incomplete archiving of flight recorder information. Check PMEM and trace files in the `/intflash/PMEM` and `/intflash/flrec` folders for all slots.

2. Display core information:

```
show core-files
```

If the timestamp for an entry in the command output matches the time the outage first occurred, transfer the core files to an FTP server. Core files are stored on the internal flash at `/intflash/coreFiles/`.

*** Note:**

Some exceptions can result in incomplete core file archiving. If the archive file list is missing files, then the system may incompletely archive the core files. If you suspect this is the case, check the staging folder at `/intflash/archive_temp` for leftover files. The `archive.sh` file in the core file archive stores the commands that the system executes when it creates the archive. The last file present in the archive generally indicates where the process stopped and the subsequent files should be looked for on the system.

Some exceptions involve reset of the master and standby CPs. Look at the capture of log files and core files on both CP file systems.

3. Obtain the network diagram of the relevant nodes, down to the port level.

Data collection for non outage problems

Use the information in this section to collect data for problems that are less service-impacting than an outage.

Gathering critical information

This section identifies the critical information that you must gather before you contact Avaya Technical Support.

You must attempt to resolve the problem using this document. Contact Avaya as a final step taken only after you are unable to resolve the issue using the information and steps provided in this document.

Gather the following information before you contact Avaya Technical Support:

- a detailed description of the problem
- the date and time when the problem started
- the frequency of the problem
- if this is a new installation
- if information exists in the InSite Knowledge Base — Have you searched the InSite Knowledge Base? Were related problem solutions found? Is there currently a work around for this issue? You can search the InSite Knowledge Base on the Avaya Support site at www.avaya.com/support. Use the Advanced Search option to narrow your search to specific categories (products) and document types.
- if the system was recently upgraded — Have you recently changed or upgraded the system, the network, or a custom application? (For example, has configuration or code been changed?) When were these changes made? Provide the date and time. Who made these changes? Were the changes made by a partner or customer? Provide the names of the individuals who made the changes.

Collecting data

Perform this procedure to collect data for problems that do not require you to restart the chassis, CP modules, or interface modules.

Procedure

1. Capture Flight Recorder trace information for the affected slot, while the problem is occurring. The flight recorder data does not provide useful information if you capture it when the problem is not occurring. If it is unclear which slot is the affected slot, then data for all slots may be necessary:

```
flight-recorder all {slot[-slot][,...]}
```

This command executes three separate commands:

- `flight-recorder snapshot`
- `flight-recorder trace`
- `flight-recorder archive`

2. Capture the current state of the chassis:

```
terminal more disable
show fulltech
```

3. Use FTP to transfer the following information:

- Configuration files from each chassis — Configuration files are stored on the internal flash at `/intflash/`.
- Log files from each chassis — Log files are stored on the external flash at `/extflash/`. If external flash does not exist, the system raises an alarm, and then logs to internal flash instead. The log file is named using the format `log.xxxxxxxx.sss` and the alarm log is named `alarmLog`.
- Generated archive files for each slot — The archive files are stored on the external flash at `/extflash/archive/[slot#]`. If the external flash does not exist, the files are stored on the internal flash at `/intflash/archive/[slot]`.

4. Show core information:

```
show core-files all
```

If the timestamp for an entry in the command output matches the time the problem first occurred, transfer the core files to an FTP server. Core files are stored on the internal flash at `/intflash/coreFiles/`.

Note:

Some exceptions can result in incomplete core file archiving. If the archive file list is missing files, then the system may incompletely archive the core files. If you suspect this is the case, check the staging folder at `/intflash/archive_temp` for leftover files. The `archive.sh` file in the core file archive stores the commands that the system executes when it creates the archive. The last file present in the archive generally indicates where the process stopped and the subsequent files should be looked for on the system.

Some exceptions involve reset of the master and standby CPs. Look at the capture of log files and core files on both CP file systems.

5. Obtain the network diagram of the relevant nodes down to the port level.

Example

The following example shows output of the `flight-recorder all` command for slot 1 only. You must use this command for all active slots as identified in the procedure steps.

```
VSP-9012:1#flight-recorder all 1
Processing Flight-recorder snapshot for 1 ....

Flight-recorder snapshot for slot 1 complete, filename is /intflash/PMEM/1/pmem.
20111019114431.1.bin.gz.

Processing Flight-recorder trace for 1 ....

Flight-recorder trace for slot 1 complete, filename is /intflash/flrec/1/trace.2
0111019114434.1.txt.

Processing Flight-recorder archive for slot 1 ....

Flight-recorder archive for slot 1 complete, filename is /intflash/archive/1/arc
hive.20111019114446.1.tar.
```

Display all show commands:

```
VSP-9012:1(config)#show fulltech
*****
Show Fulltech
Time           : Thu Jul 11 20:29:33 2013 EDT
*****

Command:[1] [ show access-policy by-mac ]
-----
default-action : allow
MAC Address    Action
=====

Command:[2] [ show access-policy snmp-group ]
-----
snmpv3-groups :
Policy 1 snmpv3-groups:
Group Name      Snmp-Model
Policy 2 snmpv3-groups:
Group Name      Snmp-Model
```

```
Command:[3] [ show access-policy ]
```

```
-----
AccessPolicyEnable: off

      Id: 1
      Name: default
PolicyEnable: true
      Mode: allow
      Service: ftp|http|telnet|ssh
      Precedence: 128
      NetAddrType: any
      NetAddr: N/A
      NetMask: N/A
      TrustedHostAddr: N/A
TrustedHostUserName: none
      AccessLevel: readOnly
      AccessStrict: false
      Usage: 0

      Id: 2
      Name: snmpv3
PolicyEnable: true
      Mode: allow
      Service: snmpv3
      Precedence: 10
      NetAddrType: any
      NetAddr: N/A
      NetMask: N/A
      TrustedHostAddr: N/A
TrustedHostUserName:
      AccessLevel: readOnly
      AccessStrict: false
      Usage: 0
```

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

Display core files information:

```
VSP-9012:1(config)#show core-files all
```

```
=====
                          Core Files
=====
Directory: /intflash/coreFiles/1
1.  File:    core.1353113115.lifecycle.CP.1.gz
    Size:    139406 bytes
    Created: Fri Nov 16 19:45:15 2012
2.  File:    core.cbcp-main.x.20121114043335.1.tar
    Size:    14059520 bytes
    Created: Wed Nov 14 04:35:36 2012
3.  File:    core.cbcp-main.x.20121114045202.1.tar
    Size:    12809728 bytes
    Created: Wed Nov 14 04:54:03 2012
4.  File:    core.cbcp-main.x.20121114050825.1.tar
    Size:    12638720 bytes
    Created: Wed Nov 14 05:10:26 2012
5.  File:    core.cbcp-main.x.20121114122506.1.tar
    Size:    13020160 bytes
    Created: Wed Nov 14 12:27:07 2012
6.  File:    core.1353336274.lifecycle.CP.1.gz
    Size:    139390 bytes
    Created: Mon Nov 19 09:44:34 2012
```

Data collection required for Technical Support cases

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

Chapter 4: Troubleshooting fundamentals

Use the following information to troubleshoot problems on the network.

Troubleshooting planning fundamentals

You can better troubleshoot the problems on the network by planning for these events in advance. To do this, you must know the following:

- that the system is properly installed and routinely maintained
- the configuration of the network
- the normal behavior of the network

Proper installation and routine maintenance

To prevent problems, follow proper maintenance and installation procedures. The following table lists the documents that provide maintenance and installation procedures.

Table 1: Maintenance and installation documentation

Subject area	Document
Chassis installation, environmental requirements	<i>Installing the Avaya Virtual Services Platform 9000</i> , NN46250-304
Control Processor, Switch Fabric, and interface module installation and replacement, cable routing	<i>Installing Modules in Avaya Virtual Services Platform 9000</i> , NN46250-301
Cooling module installation and removal	<i>Installing Cooling Modules in Avaya Virtual Services Platform 9000</i> , NN46250-302
Optical component installation and cleaning	<i>Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000</i> , NN46250-305
Power supply installation and removal	<i>Installing AC Power Supplies in Avaya Virtual Services Platform 9000</i> , NN46250-303

Network configuration

To keep track of the network configuration, gather the information described in the following sections. This information, when kept up-to-date, is extremely helpful for locating information if you experience network or device problems.

Site network map

A site network map identifies where each device is physically located on site, which helps locate the users and applications that a problem affects. You can use the map to systematically search each part of the network for problems.

Logical connections

Avaya Virtual Services Platform 9000 supports virtual LANs (VLAN). With VLANs, you must know how the devices connect logically as well as physically.

Device configuration information

Maintain online and paper copies of the device configuration information. Store all online data with the regular data backup for the site. If the site does not use a backup system, copy the information onto an external storage device, and store the backup at an offsite location.

You can use the File Transfer Protocol (FTP) and Trivial FTP (TFTP) to store configuration files on a remote server.

Other important data about the network

For a complete picture of the network, have the following information available:

- all passwords

Store passwords in a safe place. A good practice is to keep records of previous passwords in case you must restore a device to a previous software version and need to use the old password that was valid for that version.

- device inventory

Maintain a device inventory, which lists all devices and relevant information for the network. The inventory allows you to easily see the device type, IP address, ports, MAC addresses, and attached devices.

- MAC address-to-port number list

If you do not manage the hubs or switches, you must keep a list of the MAC addresses that correlate to the ports on the hubs and switches.

- change control

Maintain a change control system for all critical systems. Permanently store change control records.

- contact details

Store the details of all support contracts, support numbers, engineer details, and telephone, and fax numbers.

Normal behavior on the network

If you are familiar with the network when the network is fully operational, you can more effectively troubleshoot problems that arise. To understand the normal behavior of the network, monitor the network over a long period of time. During this time, you can see a pattern in the traffic flow, such as which devices users access most, or when peak usage times occur.

To identify problems, you can use a baseline analysis, which is an important indicator of overall network health. A baseline serves as a useful reference of network traffic during normal operation, which you can then compare to captured network traffic while you troubleshoot network problems. A baseline analysis speeds the process of isolating network problems. By running tests on a healthy network, you compile normal data for your network. You can compare this normal data against the results that you get when the network experiences trouble.

For example, ping each node to discover how long it typically takes to receive a response from devices on your network. Capture and save each response time and you can use these baseline response times to help you troubleshoot. You can also use the `show tech` and `show khi performance {buffer-pool|cpu|memory|process|pthread|slabinfo}` commands to obtain baseline output for normal system behavior.

Example

Obtain baseline output for normal system behavior:

```
VSP-9012:1#show khi performance memory
  Slot:1
    Used: 872164 (KB)
    Free: 1171940 (KB)
    Current utilization: 42 %
    5-minute average utilization: 42 %
    5-minute high water mark: 42 (08/01/11 02:09:59)

Error: Slot 2 is not active

Error: Slot 3 is not active
  Slot:4
    Used: 163588 (KB)
    Free: 320348 (KB)
    Current utilization: 33 %
    5-minute average utilization: 33 %
    5-minute high water mark: 33 (06/27/11 15:05:21)

Error: Slot 5 is not active

Error: Slot 6 is not active

Error: Slot 7 is not active

--More-- (q = quit)

VSP-9012:1#show tech

Sys Info:
-----
```

```

General Info :

  SysDescr      : VSP-9012 (4.0.0.0) (GA)
  SysName       : VSPROF
  SysUpTime    : 0 day(s), 00:32:15
  SysContact    : http://support.avaya.com/
  SysLocation   : 211 Mt. Airy Road,Basking Ridge,NJ 07920

Chassis Info:

  Chassis       : 9012
  Serial#       : LBNTMRJ0000DF
  H/W Revision  : 323500-A 01
  H/W Config    : 0100
  NumSlots      : 12
  NumPorts      : 122
  BaseMacAddr   : 80:17:7d:75:00:00
  MacAddrCapacity : 4096

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

Troubleshooting fundamentals

This section provides conceptual information and helpful tips for common problems.

Connectivity problems

Use the following general tasks to isolate connectivity problems:

- Check physical connectivity. Verify if an alarm for link or port down exists.
- Check the link state by viewing the `show interface {gigabitEthernet|loopback|mgmtEthernet|vlan}` command output.
- Use tools like ping or trace to verify if the connectivity issue is with an individual port or VLAN.
- Try to localize the affected range of ports and slot.

If you contact technical support staff to help troubleshoot connectivity problems, always provide source and destination IP pairs to facilitate in troubleshooting. Be sure to provide both working and non-working pairs for comparison.

Example

Check the link state:

```

VSP-9012:1#show interface vlan

=====
                                Vlan Basic
=====
VLAN
ID   NAME          TYPE          INST
ID   NAME          TYPE          ID   PROTOCOLID  SUBNETADDR    SUBNETMASK
=====
  
```

```

1      Default      byPort      0  none      N/A      N/A
3998  VLAN-3998    byPort      1  none      N/A      N/A
4000  RIP          byPort      1  none      N/A      N/A
    
```

All 3 out of 3 Total Num of Vlans displayed

```

=====
                                Vlan Port
=====
VLAN PORT      ACTIVE      STATIC      NOT_ALLOW
ID  MEMBER      MEMBER      MEMBER      MEMBER
-----
1   4/1-4/5,4/8-4/36,  4/1-4/5,4/8-4/36,
    4/38-4/48          4/38-4/48
--More-- (q = quit)
    
```

Routing table problems

Routing table problems include but are not limited to:

- Inactive routes
- Unnecessary routes
- Black hole routes
- Flapping links (links that go up and come down) that cause the routes to flap
- Incorrect route tables
- Invalid Address Resolution Protocol (ARP) cache that causes incorrect IP assignment
- Problems with administrative distance or other parameters

You can delete static or dynamic routes from the routing table. You can also force the device to recalculate the Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) route selection algorithms. As a last resort, you can clear the routing table and force the device to relearn routes.

Do not restart a device to clear a problem. In restarting the device, you also clear the logs. Logs are vital and can help determine many problems.

LED indications of problems

The following table lists possible problems indicated by the LEDs on Virtual Services Platform 9000 modules and suggests corrective action.

Table 2: LED problem indicators

Symptom	Probable cause	Corrective action
Green AC OK power supply LEDs are off.	The switch is not receiving AC power or the power supply has failed.	Verify that each AC power cord is fastened securely at both ends and that power is available at each AC power outlet. Plug in a device, for example, a lamp, to ensure that the power outlet is operational. Verify that each power supply is turned on.
The Link/Activity or port LED for a connected port is off or does not blink (and you believe that traffic is present).	The switch is experiencing a port connection problem, or the link partner is not auto-negotiating properly.	Verify that the cable connections to the link partner are correct. Verify port configuration parameters for both ends of the connection. Move the cable to another port to see whether the problem occurs on the new port.
The Link/Activity or port LED blinks continuously.	The switch can experience a high traffic load or possible packet broadcast storm.	Verify port configuration parameters for both ends of the connection.
The online LED is steady amber for longer than three minutes.	<p>This LED shows steady amber at module reset. This is normal behavior.</p> <p>The LED turns off before the start of the operating system, and then transitions to slow blinking amber. The LED transitions to fast blinking amber during image synchronization.</p> <p>On the IO, SF, and standby CP modules, the LED transitions to medium blinking green after the module is up and waiting for communication with the master CP module.</p> <p>On the master CP module, the LED transitions to medium blinking green waiting for communication with all the IO, SF, and standby CP modules.</p> <p>On the IO, SF, and standby CP modules, the LED transitions to steady green after communication with the master CP module establishes. On the master CP</p>	Not applicable.

Table continues...

Symptom	Probable cause	Corrective action
	module, the LED transitions to steady green after communication establishes with all other modules and the system transitions to the ready state.	
System temperature LED on master CP module is steady red.	One or more modules exceeds the normal operating temperature.	<p>Identify the module that exceeds the normal operating temperature. The online LED on the module that exceeds the normal operating temperature changes color from steady green to steady red.</p> <p>Investigate a possible cooling module failure.</p> <p>The monitoring logic polls the hardware every 30 seconds, and if the CP module reaches 55°C, the system initiates an SNMP trap and the module online LED displays a blinking and red. The system reports a temperature of 60°C as an over-temperature condition. In this case, the system automatically powers off the module. This action protects the offending module and adjacent hardware from the risk of permanent damage.</p> <p>During over-temperature conditions, the system raises an alarm and generates an SNMP trap. You can also use ACLI to obtain current temperature readings.</p>
The cooling module LED on the master CP module is steady amber.	One fan in a front cooling module has failed.	Replace the cooling module.
The cooling module LED on the master CP module is steady red.	Two or more fans in a front cooling module have failed or one or more fans in a back cooling module have failed.	Replace the cooling module.
No LEDs are lit.	A hardware failure is detected.	Turn the switch power off, and then turn it on again.

Cable connection problems

You can usually trace port connection problems to a poor cable connection or to an improper connection of the port cables at either end of the link. To remedy such problems, make sure that the cable connections are secure and that the cables connect to the correct ports at both ends of the link. If you use homemade cables, ensure that you wire the cables correctly.

1000BASE-T cables

1 gigabit per second (Gbps) ports operate using Category 5 UTP cabling only. Category 5 UTP cable is a two-pair cable. To minimize crosstalk noise, maintain the twist ratio of the cable up to the point of termination. The untwist at termination cannot exceed 0.5 in. (1.27 cm).

SFP and SFP+ cables

Cables for the optical transceivers vary depending on the specific device type. For more information about the cable requirements for small form factor pluggable (SFP) and SFP+, see *Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000*, NN46250-305.

QSFP+ cables

Cables for the optical transceivers vary depending on the specific device type. For more information about the cable requirements for quad small form factor pluggable plus (QSFP+), see *Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000*, NN46250-305.

QSFP+ authentication error

A quad small form factor plus (QSFP+) authentication error message can occur when the system detects a duplicate serial number. If this occurs, you will see the following log message:

```
CPI [07/30/14 15:48:02.969] 0x0000c609 00000000 GlobalRouter HW ERROR Duplicate QSFP+
(Serial#: 14DE215F0122 ) detected on port 7/9. Already present in port 7/5. Shutting
port 7/9
```

```
IO4 [08/11/14 19:23:43.918] 0x00318625 00000000 GlobalRouter PORT/L1 ERROR QSFP+
authentication failed for slice port 8
```

Use the `show interfaces gigabitEthernet state {slot/port[-slot/port]}[,...]` command to view the state of the port. If the port is down because it has a serial number that already exists, the port state displays as `DUP_QSFP`. View the following output to see how the command displays.

```
VSP-switch:1(config-if)#show interfaces gigabitEthernet state 7/9
```

```
=====
                                     Port State
=====
```

PORT NUM	ADMINSTATUS	PORTSTATE	REASON	DATE
7/9	up	down	DUP_QSFP	07/30/14 15:48:02

```
=====
```

Alarm database

Avaya Virtual Services Platform 9000 contains a local alarms mechanism. Applications that run on the switch raise and clear local alarms. View active alarms by using the `show alarm database` command in ACLI. Local alarms are an automatic mechanism run by the system that do not require additional configuration.

Check local alarms regularly to ensure no alarms require additional attention. The raising and clearing of local alarms also creates a log entry for each event. For more information about viewing logs, see [Viewing logs](#) on page 72.

View the alarm database regularly to monitor alarm conditions, even if you do not observe a performance problem. Review the alarm messages to determine if the system performs as expected.

Not all alarm conditions indicate a problem so you must be familiar with expected behavior.

Example

The alarm database can show the following alarm text:

```
0x00010756 Module [value] in slot [value] is non-operational
```

This alarm means that the module in the specified slot is not operating normally. The alarm typically means that the system took the module offline for some reason. If the module specified in the alarm is a CP module in slot 1 or 2, this means that the system is no longer running with the expected level of CP redundancy. Either the system took the backup CP offline, or the master CP experienced a failure and the system switched over to the backup CP to maintain proper system operation. Review the log files to determine what caused the CP failure. Service the backup CP module to return the system to the desired level of redundancy.

If the logs show the failure is a transient problem, reapply power to the CP module with the following command:

```
sys power slot <1|2>
```

This command reapplies power to the CP module. The CP module rejoins the system after it boots normally. If the problem persists and the system takes the CP out of service once again, call Avaya Support and return the CP module for service.

Troubleshooting tool fundamentals

This section provides conceptual information about the methods and tools that you can use to troubleshoot and isolate problems in the Avaya Virtual Services Platform 9000 network.

Troubleshooting overview

The types of problems that typically occur with networks involve connectivity and performance. Virtual Services Platform 9000 supports a diverse range of network architectures and protocols, some of which maintain and monitor connectivity and isolate connectivity faults.

In addition, Virtual Services Platform 9000 supports a wide range of diagnostic tools that you can use to monitor and analyze traffic, capture and analyze data packets, trace data flows, view statistics, and manage event messages.

Avaya tailors certain protocols for troubleshooting specific Virtual Services Platform 9000 network topologies. Other tools are more general in their application and you can use them to diagnose and monitor ingress and egress traffic on Virtual Services Platform 9000.

If connectivity problems occur and the source of the problem is unknown, it is usually best to follow the Open Systems Interconnection (OSI) network architecture layers. Confirm that your physical

environment, such as the cables and module connections, operates without failures before moving up to the network and application layers.

To gather information about a problem, consider the following information:

- Consider the OSI model when you troubleshoot. Start at Layer 1 and move upwards. The Address Resolution Protocol (ARP) can cause problems; ARP operates at Layer 2 to resolve MAC addresses to IP addresses (Layer 3).
- Device-specific tools and protocols can help you gather information. This document outlines Virtual Services Platform 9000-specific tools.
- You can use client- and server-based tools from Microsoft, Linux, and UNIX. For example, you can use Windows tools like ifconfig, ipconfig, winipcfg, and route print to obtain IP information and routing tables. Servers also maintain route tables.

The following command output shows example output of the route print command.

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.
C:\Documents and Settings\jsmith>route print
=====
Interface List
0x1 ..... MS TCP Loopback interface
0x2 ...00 12 f0 74 2a 87 ..... Broadcom NetLink (TM) Gigabit Ethernet - Packet
Scheduler Miniport
0x3 ...00 14 38 08 19 c6 ..... Broadcom NetXtreme Gigabit Ethernet - Packet
Scheduler Miniport
0x4 ...44 45 53 54 42 00 ..... Avaya IPSECSHM Adapter - Packet Scheduler
Miniport
=====
Active Routes:

```

Network	Destination	Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	192.168.0.1	192.168.0.102		26
0.0.0.0	0.0.0.0	207.179.154.100	207.179.154.100		1
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1		1
192.168.0.0	255.255.255.0	192.168.0.102	192.168.0.102		25
192.168.0.0	255.255.255.0	207.179.154.100	207.179.154.100		1
192.168.0.102	255.255.255.255	127.0.0.1	127.0.0.1		25
192.168.0.255	255.255.255.255	192.168.0.102	192.168.0.102		25
198.164.27.30	255.255.255.255	192.168.0.1	192.168.0.102		1
207.179.154.0	255.255.255.0	207.179.154.100	207.179.154.100		30
207.179.154.100	255.255.255.255	127.0.0.1	127.0.0.1		30
207.179.154.255	255.255.255.255	207.179.154.100	207.179.154.100		30
224.0.0.0	240.0.0.0	192.168.0.102	192.168.0.102		25
224.0.0.0	240.0.0.0	207.179.154.100	207.179.154.100		1
255.255.255.255	255.255.255.255	192.168.0.102	192.168.0.102		1
255.255.255.255	255.255.255.255	207.179.154.100	3		1
255.255.255.255	255.255.255.255	207.179.154.100	207.179.154.10		1

```
Default Gateway:207.179.154.100
=====
Persistent Routes:  None
```

- Other network problems can give the impression that a device has a problem. For instance, problems with a Domain Name System (DNS) server, another switch, firewall, or access point can appear to be routing problems.

Digital Diagnostic Monitoring

Use Digital Diagnostic Monitoring (DDM) to monitor laser operating characteristics such as temperature, voltage, current, and power. This feature works during active laser operation without affecting data traffic. Quad small form-factor pluggable plus (QSFP+) transceivers, Small form-factor pluggable plus (SFP+) transceivers, and small form-factor pluggable (SFP) transceivers support DDM. Use the ACLI command `show pluggable-optical-modules {basic|config|detail|temperature|voltage}` to make use of DDM functionality. DDM is enabled by default.

For the `show pluggable-optical-modules basic` command, the device reports qualified optics as Avaya in the type field. The device reports non-qualified best-effort optics as a different manufacturer in the type field. Unsupported optics display as unsupported in the type field, and do not operate in the system.

DDM generates temperature warnings and alarms if the port is administratively enabled. The device only generates other DDM warnings and alarms if the link is up. The device uses an offset of 8 degrees for warning thresholds and an offset of 5 degrees for alarm thresholds to calculate the thresholds used by the monitoring code. For example, if the warning threshold is 73 C and the alarm threshold is 78 C for the part, then the monitoring code uses 65 C as the warning threshold and 73 C as the alarm threshold. DDM high temperature alarm results in port shutdown for second generation module ports only.

The device sends traps if you enable the DDM traps send feature using the `pluggable-optical-module ddm-traps-send` command. The device always generates logs no matter if the DDM traps send feature is enabled or disabled. Configure the DDM monitor interval using the `pluggable-optical-module ddm-monitor-interval <10..40>` command. The DDM monitor interval is 10–40 seconds. The default is 10 seconds.

An interface that supports DDM is a Digital Diagnostic Interface (DDI). These devices provide real-time monitoring of individual DDI QSFP+, SFP+s, and SFPs on a variety of Avaya products. The DDM software provides warnings or alarms when the temperature, voltage, laser bias current, transmitter power, or receiver power fall outside of vendor-specified thresholds during initialization.

* Note:

Digital Diagnostic Interface (DDI) module information for RxPower can output false alerts. The MSA (Multi-source Agreement) between manufacturers of QSFP+, SFP+s, and SFP devices specifies a +/- 3dB accuracy tolerance for optical power measurements.

To minimize false warnings or alarms due to this inaccuracy, the thresholds for low and high TxPower and for low RxPower are offset by this tolerance. High RxPower thresholds are not offset due to the potential for receiver saturation and damage that can result over the long-term, however, this increases the possibility of false alerts.

If high RxPower alerts occur, but the link operates normally, consider this tolerance. If the link fails to operate, consider the possibility that the optical receiver is being over-driven, and attempt to correct the condition.

For information about DDM and QSFP+, SFP+s, and SFP, see *Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000*, NN46250-305, and *Monitoring Performance on Avaya Virtual Services Platform 9000*, NN46250-701.

Example

Display Digital Diagnostic Monitoring (DDM) functionality:

```
VSP-9012:1#show pluggable-optical-modules config
=====
Pluggable Optical Module Global Configuration
=====
          ddm-monitor : enabled
ddm-monitor-interval : 10
          ddm-traps-send : enabled
          ddm-alarm-portdown : enabled
```

Port mirroring

Virtual Services Platform 9000 has a port mirroring feature that helps you monitor and analyze network traffic. Port mirroring supports both ingress (incoming traffic) and egress (outgoing traffic) port mirroring. When you enable port mirroring, the system forwards ingress or egress packets normally from the mirrored (source) port, and sends a copy of the packet to the mirroring (destination) port.

Overview

Port mirroring causes the switch to make a copy of a traffic flow and send the copy to a device for analysis. You can use port mirroring in diagnostic sniffing. You can use the mirror to view the packets in the flow without breaking the physical connection to place a packet sniffer inline. You can also use mirroring for security reasons.

You can use egress mirroring to monitor packets as they leave specified ports.

Use a network analyzer to observe and analyze packet traffic at the mirroring port. Unlike other methods that analyze packet traffic, the packet traffic is uninterrupted and packets flow normally through the mirrored port.

You can mirror to a port or list of ports, a VLAN, or a MultiLink Trunking (MLT) group. Virtual Services Platform 9000 supports one-to-many, many-to-one, and many-to-many mirroring configurations.

Port mirroring and modules

You can use all module ports in the system to function as an ingress port for mirroring (mirrored port), an egress port for mirroring (mirrored port), or as a mirroring port (where all the mirrored traffic is redirected). The number of mirroring ports (also called destination ports) that you can configure depends on the quantity of modules you have in your system configuration. The software limitation is 479 ports simultaneously.

The following table describes ingress mirroring functionality for modules. The system only supports one type of mirroring destination at a time. You cannot mirror the same port to multiple classes of destinations, for example, MLT and VLAN. However, you can mirror to multiple physical destinations.

Table 3: Ingress mirroring functionality

Function	Support information
Ingress port mirroring and ingress flow mirroring	Supported, no restriction for each lane
One port to one port	Supported, no restriction for each lane Layer 3 supports one-to-one for both port and flow-based remote mirroring
One to MLT group [for threat protection system (TPS applications)]	Supported
One to many (multicast group ID/VLAN)	Supported
One to one remote mirrored destination	Supported
Many to one (multiple mirrored ports to one mirroring port)	Supported Layer 3 supports many-to-one for both port and flow-based remote mirroring
Many to MLT group	Supported
Many to many (VLAN/multicast group ID) (multiple ports with several different destinations)	Supported
Many to one relation between Remote Mirror Source (RMS) and Remote Mirror Termination (RMT)	Supported
VLAN and port combination as a mirroring destination	Not supported
Ingress flow mirroring	Supported
Allow filters to specify a separate destination for each access control entry	Supported
Flow-based remote mirroring	Supported for Layer 3

The following table describes egress mirroring functionality.

Table 4: Egress mirroring functionality

Function	Support information
Egress port mirroring and egress flow mirroring	Supported
One port to one port	Supported Layer 3 supports one-to-one for both port and flow-based remote mirroring
One to MLT groups (for TPS applications)	Supported
One to many (multicast group ID/VLAN)	Supported
One to one remote mirrored destination	Supported
Many to one (multiple mirrored ports to one mirroring port)	Supported Layer 3 supports many-to-one for both port and flow-based remote mirroring

Table continues...

Function	Support information
Many to MLT group	Supported
Many to many (multicast group ID) (multiple ports with several different destinations)	Supported
Many to one relation between Remote Mirror Source (RMS) and Remote Mirror Termination (RMT)	Supported
VLAN and port combination as mirroring destination	Not supported
Egress flow mirroring	Supported
Allow filter to specify a separate destination for each access control entry	Supported
Flow-based remote mirroring	Supported for Layer 3

Multiport mirroring uses multicast group IDs (MGID) to perform mirroring and replicate it to all the mirrored interfaces. If multiple mirroring interfaces exist, the CP module allocates an MGID to that mirrored stream. The maximum number of system MGIDs available for port mirroring, along with flow-based mirroring, is 176. If you use the same mirroring ports for different instances of mirroring configuration, the system uses the same MGID.

Module configuration

You can specify a destination multilink trunking (MLT) group, a destination port or set of ports, or a destination VLAN.

Interface modules support two port mirroring modes: rx (ingress, which is, inPort and inVLAN) and tx (egress, which is, outPort and outVLAN). Configure the mirroring action globally in an access control list (ACL), or for a specific access control entry (ACE) by using the ACE mirror actions. Configure the mirroring destination by using an ACE.

In rx modes, when you configure the ACE mirror or ACL global options to mirror, use the ACE to configure the mirroring destination port.

To modify a port mirroring instance, first disable the instance. Also, to change a port, VLAN, or MLT entry, first remove whichever parameter is attached to the entry, and then add the required entry. For example, if an entry has mirroring ports already assigned, then you have to remove the ports using the `no mirror-by-port` command, and then, to assign a VLAN to the entry, use the `mirror-by-port monitor-vlan` command.

ACLs, ACEs, and port mirroring

You can configure an ACL or an ACE to perform the mirroring operation. To do so, you can configure the ACL global action to mirror, or you can configure the ACE action to mirror. If you use the global action, mirroring applies to all ACEs that match in an ACL.

To decouple flow-based mirrors from port-based mirrors, ACEs use a parameter called mirror, which you can configure to a specific mirror to MLT ID, VLAN, port, or port list.

You can use filters to reduce the amount of mirrored traffic. To use filters with port mirroring, you must use an ACL-based filter. Apply an ACL to the mirrored port in the egress and ingress directions. The system forwards traffic patterns that match the ACL or ACE with an action of permit to the destination and also to the mirroring port. The system does not forward traffic patterns that match an ACE with an action of drop (deny) to the destination, but those traffic patterns still reach the mirroring port. For example, for an ACL or ACE with a match action of permit and debug

mirroring enabled, the system mirrors packets to the specified mirroring destination on the ACE. If you enable a port or VLAN filter, that filter is the mirroring filter.

You can specify more than one mirroring destination by using multiple ACEs. Use each ACE to specify a different destination.

You can configure a port-based and a flow-based mirroring filter on the same port. If such a case occurs, the flow-based mirror takes precedence.

For more information about how to configure ACLs and ACEs, see *Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000*, NN46250-502.

Port mirroring considerations and restrictions

Although you can configure Virtual Services Platform 9000 to monitor both ingress and egress traffic, some restrictions apply:

- Mirrored traffic shares ingress queue and fabric bandwidth with normal traffic and therefore can impact normal traffic. Therefore, use these features with this potential consequence in mind, and enable them only for troubleshooting, debugging, or for security purposes such as packet sniffing, intrusion detection, or intrusion prevention.
- You can configure as many ingress mirroring flows as you have filters.
- To avoid VLAN members from seeing mirrored traffic, you must remove mirroring (destination) ports from all VLANs.
- The MAC drops an errored packet, for example, packets that are too short or too long. Control packets consumed by the MAC (802.3x flow control) are also not mirrored.
- You cannot use port mirroring on operations, administration, and maintenance (OAM) ports.

Remote mirroring

Use remote mirroring to steer mirrored traffic through a switch cloud to a network analysis probe located on a remote switch. Use remote mirroring to monitor many ports from different systems by using one network probe device and encapsulating mirrored packets.

Layer 2 remote mirroring

The encapsulated frame can be bridged through the network to the remote diagnostic termination port.

Remote mirroring uses a specific VLAN if you enable remote mirroring on the port mirroring destination port. The VLAN ID is in the monitor tag field of the remote mirrored packet. With this feature, you can segregate remote mirrored traffic to a single VLAN.

In addition, you can monitor traffic for Media Access Control (MAC) addresses, where the system copies traffic with a certain MAC source address (SA) or MAC destination address (DA) to the specified mirroring port. You can use the VLAN forwarding database feature to monitor traffic for Media Access Control (MAC) addresses. In this case, the system copies traffic with a certain source or destination MAC address to the mirror port. Monitoring of MAC address traffic must be within the context of a VLAN.

When an RMT port receives an encapsulated frame from the switch fabric, it strips off the remote mirroring encapsulation as the system transmits it on the port. The system identifies remote mirrored encapsulated frames when the system detects the configured remote mirroring destination MAC

address as the destination MAC address in the packet. The RMT sends dummy broadcast Layer 2 packets with the remote mirroring destination MAC address as the source MAC address so that all nodes in the network can learn this MAC address. The RMT sends this broadcast every 10 seconds because the minimum value of the forwarding database (FDB) aging timer is 10 seconds. After you configure a port as an RMT, the system adds a static FDB entry to channel all traffic destined for the remote mirroring destination MAC address to the RMT port. When you remove an RMT port from all of the configured VLANs, the system disables the remote mirroring feature on the port.

The remote mirroring encapsulation wrapper is 20 bytes in length and consists of a Layer 2 destination address, Layer 2 source address, monitor tag, monitor Ethertype, and monitor control. The system strips the original CRC-32 from a mirrored packet, and computes a new CRC-32 over the entire encapsulated frame. When the mirrored frame is 1522 bytes (1518 plus 4-byte 802.1p/q tag), the resulting maximum frame length is 1542 bytes. To support this, all the nodes in the network must be able to handle 1542-byte packets.

The following figure illustrates Layer 2 remote mirroring with four Virtual Services Platform 9000 systems and a client with a network analysis probe.

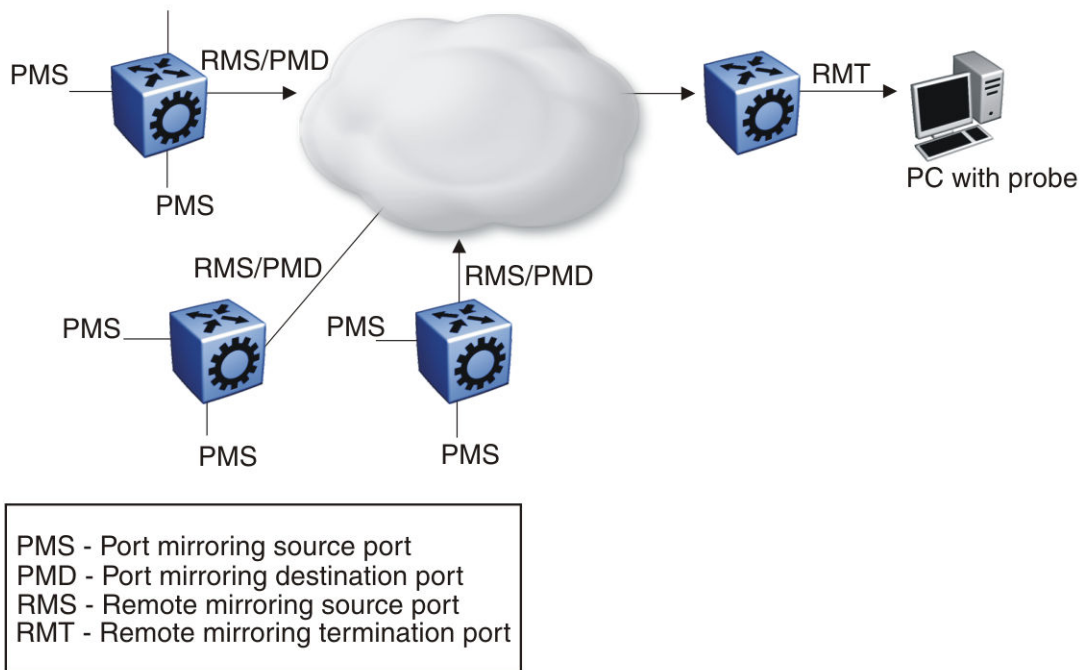


Figure 1: Layer 2 remote mirroring

Layer 2 remote mirroring considerations and restrictions:

Mirrored traffic shares ingress, egress, and fabric bandwidth with normal traffic and can impact normal traffic. Use these features with this potential consequence in mind and enable them only for troubleshooting, debugging, or for security purposes, such as packet sniffing, intrusion detection, or intrusion prevention.

To support remote mirroring, all the nodes in the network must be able to handle a packet size of up to 1542 bytes.

You can create multiple remote mirroring source (RMS) or remote mirroring termination (RMT) ports in each lane on a module.

The following limitations apply to remote mirroring:

- The system supports a maximum of 32 RMT ports.
- The RMS port must be a port mirroring destination port because only mirrored packets are remote mirrored. The platform does not check if the port is a port mirroring destination port, and sends no error messages if the port is not.
- An RMT must be part of at least one port-based VLAN.
- If a port mirroring entry exists with remote mirroring enabled on a particular VLAN, you cannot convert the VLAN to a routable VLAN.
- You cannot use remote mirroring on operations, administration, and maintenance (OAM) ports.

Note the following information:

- If the RMS is a tagged port, the system encapsulates and transmits the mirrored packet with the VLAN ID of the RMS port and forwards to the RMT. Encapsulation does not modify the mirrored packet data or the VLAN ID. When the RMT port receives an encapsulated frame from the switch fabric, the port removes the remote mirroring encapsulation and the system transmits the frame on the port with the VLAN ID of the mirrored packet (the original packet).
- If you disable port mirroring, the system does not remote mirror any packets.
- The system captures packets as long as the RMT is reachable.
- When you enable or disable Layer 2 remote mirroring, the system sends a trap to the trap receiver, and an SNMP log message states that remote mirroring is enabled or disabled, and the mode.
- For Layer 2 remote mirroring, when you remove an I/O module from a slot, the RMS and the RMT on all ports in the slot are disabled. This action generates an SNMP log message and a trap. When you reinsert the module, the RMS and RMT are reenabled along with remote mirroring.

Layer 3 remote mirroring

Virtual Services Platform 9000 supports Layer 3 remote mirroring for ports and flows. Use Layer 3 mirroring to monitor traffic remotely. Layer 3 remote mirroring monitors traffic from multiple network devices across an IP network, and sends that traffic in an encapsulated form to the destination analyzers.

Specify the destination as an IP address. The source of the encapsulated packet and destination interfaces can be on different devices connected by an IP network. The Layer 3 remote mirror traffic is Global Routing Engine (GRE) encapsulated. Encapsulated ports can be MLT or VLAN ports. The following figure shows a Layer 3 remote mirroring configuration.




Figure 2: Layer 3 remote mirroring

In the preceding figure, a network analyzer monitors the ingress and egress traffic for VSP-A using GRE encapsulation. Encapsulated packets are routed from VSP-A through the routed network to the destination device (VSP-B), which decapsulates the packets and forwards them to the attached network analyzer.

Layer 3 remote mirroring supports the following configurations for both port- and flow-based mirroring: remote mirroring and flow mirroring

- One-to-one mirroring—This configuration supports one mirrored port and one monitored IP address.
- Many-to-one mirroring—This configuration supports multiple mirrored ports to one monitored IP address.

Virtual Services Platform 9000 supports Layer 3 remote mirroring if the learned port is an MLT or VLAN port. The route can be a default route, default ECMP route, ECMP route, or dynamically learned.

For Layer 3 remote mirroring, every hop in the path from the mirrored port to the remote-mirroring port must be routed. Avaya recommends that you configure the remote-mirrored port in its own VLAN at the last hop to prevent flooding.

Layer 3 remote mirroring considerations and restrictions:

The following limitations apply to remote mirroring:

- If a port mirroring entry exists with remote mirroring enabled on a particular VLAN, you cannot convert the VLAN to a routable VLAN.
- Virtual Services Platform 9000 does not support Layer 3 remote mirroring for multiple IP destinations.
- If the end device is not a Virtual Services Platform 9000, the GRE encapsulated packet is routed to the monitor destination on the end device.
- If the end device is not a Virtual Services Platform 9000 and has bridging to the monitor destination, you must enable port mirroring on the end device to see the source packets transmit to the monitor destination.
- If the monitor destination is on another Virtual Services Platform 9000 and the system drops packets, you can view this information by using the `show khi forwarding rsp` command. The number of packets appears under `L3MirrorDrops`.

- If the end device is a Virtual Services Platform 9000 and it uses bridging to the monitor destination, the mirror packets are dropped before RSP. They are dropped at the MAC level.
- You cannot configure an ACL global action for a Layer 3 mirror.
- Avoid zero IP addresses, broadcast addresses, loop back addresses, and other invalid addresses.
- Do not use the IP address of the master or backup CP module.
- Do not monitor a virtual management IP address.
- Verify that the optional DSCP and TTL parameters use valid ranges.
- Do not monitor the remote VLAN ID.
- Do not configure monitor-ip on the same subnet as an interface on the chassis. If a mirror entry exists with monitor-ip and you configure an IP address, which is in the same subnet as monitor-ip, the system restricts the IP address creation based on the mirror configuration.
- You cannot use remote mirroring on operations, administration, and maintenance (OAM) ports.
- Layer 3 remote mirroring may not work when monitoring destinations on SMLT VLANs. Users should have the Layer 3 mirroring destination connected on a separate VLAN. This VLAN should not span on multiple switches.
- Layer 3 flow mirroring does not take precedence over port mirroring on second generation I/O modules. The mirror destination will receive both copies from port and flow mirroring.

Note the following information:

- If you disable port mirroring, the system does not remote mirror any packets.
- Layer 3 remote mirroring supports trace; it does not support log messages and traps.

Packet Capture Tool

The Packet Capture Tool (PCAP) is a data packet capture tool that captures ingress and egress packets on selected I/O ports. With this tool, you can capture, save, and download one or more traffic flows through Virtual Services Platform 9000. You can then analyze the captured packets offline for troubleshooting purposes. This tool uses the mirroring capabilities of the I/O ports.

To use PCAP, you must have the Advanced Software License. For more information about licensing, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

The secondary CP module acts as the PCAP engine and stores all captured packets. The master CP module maintains protocol handling and capture activity does not affect the master CP module.

PCAP packet flow

By default, PCAP uses port mirroring. If you apply a filter set, PCAP uses flow mirroring. If you require further filtering, apply PCAP software filters. You can store captured packets in the PCAP engine DRAM (PCAP00) or on the network. You can then use the File Transfer Protocol (FTP) to download the stored packets to an offline analyzer tool such as EtherReal or Sniffer Pro.

The following figure illustrates how to use the PCAP tool to configure PCAP filters and enable them on ports.




Figure 3: PCAP configuration

PCAP feature support

PCAP supports the following features:

- PCAP uses the secondary CP module as the PCAP engine.
- PCAP supports activating packet capture on one or multiple ports.
- PCAP can capture packets on ingress, egress, or both directions.
- PCAP supports software filters, which provide a way to filter the packets in the PCAP engine.
- Captured packets can be stored on a Compact Flash device or on the network. The packets are stored in Sniffer Pro file format.

PCAP filters

Use the PCAP filters to selectively configure match criteria to capture or drop frames. The configured parameters determine which filter to apply to a frame. The default behavior is to accept the frame. You can also configure trigger filters to globally start and stop packet capturing.

If you enable PCAP using capture filters with the action `trigger-on`, after the first packet that matches the filter criteria hits the PCAP engine, the system disables capture filter and PCAP capture starts. If you enable PCAP using capture filters with the action `trigger-off`, PCAP captures all

packets until the first one that matches the filter criteria, and then disables the capture filter and globally disables PCAP.

Because the PCAP engine runs on the secondary CP module, the master CP module does not reflect the change in PCAP and PCAP capture-filter status if you use the action **trigger-on** or **trigger-off**. Run the **show pcap capture-filter** or **show pcap cli** commands on the secondary CP module to view the correct status. After PCAP disables the filter entry on the PCAP engine, if you use the **show pcap** or **show pcap capture-filter** command on the master CP module, the status appears as true (enabled), when it really is false (disabled). To activate PCAP and the PCAP capture filter again, you must reenable them on the master CP module.

The following example shows the status line in the command output on the secondary CP module.

```
VSP-9012:2#show pcap
enable = FALSE
buffer-wrap = TRUE
wrap-auto-save-file = TRUE
buffer-size = 32 MB
fragment-size = 64 Bytes
auto-save = TRUE
AutoSaveFilename = pcap.cap
AutoSaveDevice = extflash

VSP-9012:2#show pcap capture-filter

=====
PCAP Capture-filters
=====

Id: 1
action : trigger-on
enable : false
srcmac : 00:00:00:00:00:00 Mask = 6
dstmac : 00:00:00:00:00:0a Mask =6
srcip : 0.0.0.0 to 0.0.0.0
dstip : 0.0.0.0 to 0.0.0.0
vlan-id : 0 to 0
pbits : 0 to 0
ether-type : 0x0 to 0x0
protocol-type : 0 to 0
dscp : 0 to 0
udp-port : 0 to 0
tcp-port : 0 to 0
user-defined: Offset: 0 Data:
timer : 1000 ms
packet-count : 0
refresh-timer : 0 ms
```

The following table explains how to use the capture filter to achieve the desired results.

Table 5: Capture filter examples

Example	PCAP configuration	Resulting action
PCAP capture filter with action trigger-off, match dstmac, and interface mode rx.	interface gigabitEthernet 9/11 pcap enable mode rx exit pcap capture-filter 1	This capture filter captures packets in the PCAP engine until it receives the first packet with destination MAC 0x00:00:00:00:00:0a. After the

Table continues...

Example	PCAP configuration	Resulting action
	<pre>pcap capture-filter 1 dstmac 00:00:00:00:00:0a pcap capture-filter 1 action trigger-off pcap capture-filter 1 enable pcap enable</pre>	<p>engine receives the matching packet, the system disables the capture filter and PCAP globally. The PCAP engine does not capture more packets.</p>
<p>PCAP capture filter with action trigger-on, match dstmac, and interface mode tx.</p>	<pre>interface gigabitEthernet 6/5 pcap enable mode tx exit pcap capture-filter 1 pcap capture-filter 1 dstmac 00:00:00:00:00:0a pcap capture-filter 1 action trigger-on pcap capture-filter 1 enable pcap enable</pre>	<p>This capture filter drops packets in the PCAP engine until it receives the first packet with destination MAC 0x00:00:00:00:00:0a. After the engine receives the matching packet, the system disables capture filter and the engine captures all packets, starting with the matched one, until you disable PCAP manually.</p>
<p>PCAP capture filter with action capture, match dstmac, and interface mode both.</p>	<pre>interface gigabitEthernet 9/11 pcap enable mode both exit pcap capture-filter 1 pcap capture-filter 1 action capture pcap capture-filter 1 dstmac 00:00:00:00:00:0a pcap capture-filter 1 enable pcap capture-filter 2 pcap capture-filter 2 enable pcap capture-filter 2 action drop pcap enable</pre>	<p>The configured filters capture packets with destination MAC 0x00:00:00:00:00:0a and drop the rest from the PCAP engine. You must use the second filter to ensure that PCAP drops all packets that do not match the capture filters.</p>
<p>PCAP capture filter with action trigger-on with timer, match dstmac, and interface mode rx.</p>	<pre>interface gigabitEthernet 9/11 pcap enable mode rx exit pcap capture-filter 1 pcap capture-filter 1 action trigger-on pcap capture-filter 1 dstmac 00:00:00:00:00:0a pcap capture-filter 1 timer 1000</pre>	<p>This configuration enables a timer. After the PCAP engine receives a packet with destination MAC 0x00:00:00:00:00:0a, it captures all packets for the duration of the timer, and then disables PCAP globally. Specify the timer value in milliseconds (ms). A value of 1000 equals 1 second.</p>

Table continues...

Example	PCAP configuration	Resulting action
	<pre>pcap capture-filter 1 enable pcap enable</pre>	
<p>PCAP capture filter with action trigger-on with refresh-timer, match dstmac, and interface mode rx.</p>	<pre>interface gigabitEthernet 9/11 pcap enable mode rx exit pcap capture-filter 1 pcap capture-filter 1 action trigger-on pcap capture-filter 1 dstmac 00:00:00:00:00:0a pcap capture-filter 1 refresh- timer 60000 pcap capture-filter 1 enable pcap enable</pre>	<p>This configuration enables a timer. After the PCAP engine receives a packet with destination MAC 0x00:00:00:00:00:0a, it disables the capture filter. After the PCAP engine receives the matching packet, and if it does not receive more matching packets for the duration of the timer, it disables PCAP globally.</p> <p>The timer restarts every time the PCAP engine receives a packet, until the timer expires. Specify the timer value in ms. A value of 60000 equals 1 minute.</p>

PCAP limitations and considerations

This section describes the limitations and considerations of the PCAP tool.

- Flow control packets can be issued if port performance is affected while PCAP is enabled.
- When you configure capture-filter parameters for PCAP, the software accepts a value of 0 for the range of values. The value of 0 disables the filter parameter. Do not use 0 in a range of values in a filter parameter.
- When the secondary CPU cycles in the PCAP engine are used for packet capturing, and if the packet incoming rate is high (approximately 200 Mb/s), the log messages and certain commands executed in the secondary CPU are queued until the packet capturing is complete. For immediate recovery, disable PCAP on the individual ports in the master CPU on which packets are to ingress. The packets captured are stored in the buffer.
- To autosave by using an anonymous FTP session to a Windows system, first create a /pub subdirectory in the c: drive or the drive that is the default for the FTP server.
- PCAP uses two levels of filtering to capture packets: one at the hardware level and one at the software level. The hardware level uses PCAP filters; the software level uses capture filters. Therefore, when you use the `show pcap port` command, you can see filter set values that are specific to IP traffic filters only.

Use the `pcap enable` command to enable or disable PCAP on the port. When you use the `show pcap port` command, the information that appears refers to PCAP only (If enable is configured to true, this means that PCAP is enabled for the specified interface).

- If you use a PCAP filter to capture packets, and then you disable PCAP globally and at the port level, the filter remains active.
- If you globally disable PCAP, the number of packets dropped in hardware continues to go up unless you also disable PCAP on the port. To disable PCAP on the port, use the `no pcap enable` command.

- If the chassis is in HA mode, after a PCAP buffer wrap occurs, a log message appears on the console of the primary CPU to indicate that the buffer has wrapped. If the chassis is not in HA mode, the log message appears only on the secondary CPU. This difference exists because the CPUs do not synchronize log messages in non-HA mode.
- You cannot use PCAP on operations, administration, and maintenance (OAM) ports.

PCAP and I/O modules

At the port level, you can enable PCAP in one of the following modes:

- rx (ingress)
- tx (egress)
- both (both ingress and egress)

Flight Recorder

The Flight Recorder is a high level term for the framework in place on Virtual Services Platform 9000 to store both history and current state information for various kernel, system, and application data with minimal overhead to execution. You can later access this data on-demand when you debug system issues to give engineers the best possible troubleshooting information. Functionally, the Flight Recorder consists of two elements: Persistent Memory and Always-on Trace.

The Persistent Memory feature stores information in volatile memory that persists across processor resets. This feature provides information on crashes, errors, and outages that are not the result of a power failure. Persistent Memory data not saved to non-volatile storage before a power failure is lost. The system takes Persistent Memory snapshots when:

- A critical process stops functioning.
- A card resets.
- The hardware watchdog activates.
- The user initiates a snapshot in the ACLI.

The Always-on Trace feature creates an ongoing, circular log of every trace call recently executed regardless of the trace level enabled by the user. This functionality provides 128K of storage for control processor trace records, 32K of storage for input/output trace records, and 16K for switch fabric trace records. Since the Always-On Trace performs circular logging, reading the log from top to bottom does not represent a chronological sequence of events. Pay attention to timestamp information to discern the chronology of events.

Flight Recorder functionality is provided only through ACLI. Use the following commands to use this feature:

- **flight-recorder snapshot <slot>**

This command outputs the Persistent Memory information for the specified slot to a file. The system notifies the user of the name and location of the file at the end of the output process.

- **flight-recorder trace <slot>**

This command outputs the Always-on Trace information for the specified slot to a file. The system notifies the user of the name and location of the file at the end of the output process.

- `flight-recorder all <slot>`

The command outputs both the Persistent Memory and Always-on Trace information for the specified slot. The system notifies the user of the name and location of the files created at the end of the output process.

- `flight-recorder archive <slot>`

This command outputs a Flight Recorder archive for the specified slot to a file. The system notifies the user of the name and location of the file at the end of the output process. This archive file provides a complete set of debug information the user can provide for technical assistance.

General diagnostic tools

Virtual Services Platform 9000 has diagnostic features available with Enterprise Device Manager (EDM) and Avaya command line interface (ACLI). You can use these diagnostic tools to help you troubleshoot operational and configuration issues. You can perform such tasks as configuring and displaying log files, viewing and monitoring port statistics, tracing a route, running loopback and ping tests, testing the switch fabric, and viewing the address resolution table.

For more information about statistics, see *Monitoring Performance on Avaya Virtual Services Platform 9000*, NN46250-701.

Traceroute

Traceroute determines the path a packet takes to reach a destination by returning the sequence of hops (IP addresses) the packet traverses.

According to RFC1393, traceroute operates by: "sending out a packet with a time-to-live (TTL) of 1. The first hop then sends back an ICMP error message indicating that the packet cannot forward because the TTL expired. The packet is then resent with a TTL of 2, and the second hop returns the TTL expired. This process continues until the destination is reached. The purpose behind this is to record the source of each ICMP TTL exceeded message to provide a trace of the path the packet took to reach the destination."

Ping

Ping is a simple and useful diagnostic tool to determine reachability. When you use ping, the switch sends an ICMP echo request to a destination IP address. If the destination receives the packet, it responds with an ICMP echo response.

If a ping test is successful, the destination is alive and reachable. Even if a router is reachable, it can have improperly working interfaces or corrupted routing tables.

Trace

Use trace commands to provide detailed data collection about software modules on Virtual Services Platform 9000. You can use the trace toolset to trace multiple modules simultaneously and provide options to specify the verbosity level of the output.

You can enable trace logging through the `boot config trace-logging flag`. This command causes the trace output to be captured in systrace files in the external flash of the primary CP module. A trace run with this flag set to true is copied to the CF under file systrace.

 **Caution:**

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the switch, loss of protocols, and service degradation.

While these occurrences are uncommon, when using the trace level tool, minimize this risk. Avaya recommends:

- In situations where you require trace data concurrently from multiple modules, consider troubleshooting during a maintenance window if feasible. Consider a maintenance window period if the switch is stable but CPU utilization is high and CPU traces (example trace levels 9 and 11) are required to diagnose the cause.
- To avoid potential issues due to logging trace data to the CF card, disable the trace-logging feature (`no boot config flags trace-logging`).
- Run trace commands from the console port when the CPU utilization is already high. While you can enable or disable tracing when directly connected to the console port, Avaya recommends that you use an SSH or Telnet connection to the management port.
- Activate tracing on one software module at a time.
- Initially activate tracing at lower verbosity settings (that is, 2 rather than 3). Increase to verbosity level 3 or 4 only if required, and after level 2 runs safely.
- Avoid leaving traces active for extended periods of time. For high CPU utilizations, a few seconds (typically less than 5 seconds) is generally sufficient to identify the cause for sustained high CPU utilization.

Chapter 5: Log and trap information

Use the following information to understand Simple Network Management Protocol (SNMP) traps and log files.

Log and trap fundamentals

Use the information in this section to help you understand Simple Network Management Protocol (SNMP) traps and log files, available as part of Avaya Virtual Services Platform 9000 System Messaging Platform.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) provides facilities to manage and monitor network resources. SNMP consists of:

- Agents—An agent is software that runs on a device that maintains information about device configuration and current state in a database.
- Managers—An SNMP manager is an application that contacts an SNMP agent to query or modify the agent database.
- SNMP protocol—SNMP is the application-layer protocol used by SNMP agents and managers to send and receive data.
- Management Information Bases (MIB)—The MIB is a text file that specifies the managed objects by an object identifier (OID).

Important:

Virtual Services Platform 9000 does not reply to SNMP requests sent to the Virtual Router Redundancy Protocol (VRRP) virtual interface address; it does, however, reply to SNMP requests sent to the physical IP address.

An SNMP manager and agent communicate through the SNMP protocol. A manager sends queries and an agent responds; however, an agent initiates traps. Several types of packets transmit between SNMP managers and agents:

- Get request—This message requests the values of one or more objects.
- Get next request—This message requests the value of the next object.

- Set request—This message requests to modify the value of one or more objects.
- Get response—An SNMP agent sends this message in response to a get request, get next request, or set request message.
- Trap—An SNMP trap is a notification triggered by events at the agent.

Overview of traps and logs

SNMP traps

The SNMP trap is an industry-standard method used to manage events. You can set SNMP traps for specific types of log message (for example, warning or fatal), from specific applications, and sends them to a trap server for further processing. For example, you can configure Virtual Services Platform 9000 to send SNMP traps to a server after a port is unplugged or if a power supply fails.

This document only describes SNMP commands related to traps. For more information about how to configure SNMP community strings and related topics, see *Configuring Security on Avaya Virtual Services Platform 9000*, NN46250-601.

System log messaging

On a UNIX-based management platform, you can use system log (syslog) messaging to manage event messages. Virtual Services Platform 9000 syslog software communicates with a server software component named syslogd on the management workstation.

The UNIX daemon syslogd is a software component that receives and locally logs, displays, prints, and forwards messages that originate from sources internal and external to the workstation. For example, syslogd on a UNIX workstation concurrently handles messages received from applications that run on the workstation, as well as messages received from Virtual Services Platform 9000 that run in a network accessible to the workstation.

The remote UNIX management workstation performs the following actions:

- Receives system log messages from Virtual Services Platform 9000.
- Examines the severity code in each message.
- Uses the severity code to determine appropriate system handling for each message.

Log consolidation

Virtual Services Platform 9000 generates a system log file and can forward that file to a syslog server for remote viewing, storage and analyzing.

The system log captures messages for the following components:

- Simple Network Management Protocol (SNMP)
- Extensible Authentication Protocol (EAP)
- Remote Authentication Dial-in User Service (RADIUS)
- Remote Monitoring (RMON)
- Web
- Internet Group Management Protocol (IGMP)
- hardware (HW)

- MultiLink Trunking (MLT)
- filter
- Quality of Service (QoS)
- Command line interface (CLI) log
- software (SW)
- Central Processing Unit (CPU)
- Internet Protocol (IP)
- Virtual Local Area Network (VLAN)
- Internet Protocol Multicast (IPMC)
- Internet Protocol-Routing Information Protocol (IP-RIP)
- Open Shortest Path First (OSPF)
- policy
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP) log

Avaya Virtual Services Platform 9000 can send information in the system log file, including ACLI command log and the SNMP operation log, to a syslog server.

View logs for the CLILOG module to track all ACLI commands executed and for fault management purposes. The system logs the ACLI commands to the system log file as CLILOG module.

View logs for the SNMPLOG module to track SNMP logs. The system logs the SNMP operation log to the system log file as SNMPLOG module.

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs ACLI log and SNMP log information regardless of the logging level you set. This is not the case for other INFO messages.

System log client over IPv6 transport

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration except in the following case. When you configure the system log table, under the System Log Table tab, you must select either IPv4 or IPv6.

Log message format

The log messages for Virtual Services Platform 9000 have a standardized format. The device tags all system messages with the following information, except that alarm type and alarm status apply to alarm messages only:

- Avaya proprietary (AP) format—Provides encrypted information for debugging purposes.
- Module—Identifies the software module or hardware from which the log is generated.
- Timestamp—Records the date and time at which the event occurred. The format is MM/DD/YY hh:mm:ss.uuu, where uuu is milliseconds. Example: [11/01/10 11:41:21.376].

- Event code—Precisely identifies the event reported.
- Event instance or alarm ID—Identified the instance of the event or alarm ID for alarm messages.
- Alarm type—identifies the alarm type (Dynamic or Persistent) for alarm messages.
- Alarm status—identifies the alarm status (set or clear) for alarm messages.
- VRF name—identifies the Virtual Routing and Forwarding (VRF) instance, if applicable.
- Severity level—Identifies the severity of the message.
- Terse message—Represents the event and provides additional information.
- Probable cause—describes the possible conditions that trigger the event.

The following messages are examples of an informational message, warning message, and alarm messages:

```
IO5 [08/17/11 11:38:04.875] 0x0009059e 00000000 GlobalRouter QOS INFO QOS profile set to 0
SF4 [08/17/11 11:38:04.875] 0x0009059e 00000000 GlobalRouter QOS INFO QOS profile set to 0
CP1 [08/16/11 11:38:04.875] 0x00043fff 00000000 GlobalRouter WEB INFO HTTPS: Server Cert/
Key Generated Successfully
```

Unprintable characters in log and trace messages are encoded in C language hexadecimal notation in the string and surrounded with the <UNPRINTABLE>...<UNPRINTABLE> tag. For example, the device converts a 4-byte sequence of 0x01, 0x02, 0xfe, 0xff to \x01\x02\xfe\xff, and adds the <UNPRINTABLE>...<UNPRINTABLE> tag to the log or trace message. The following example displays how the message appears:

```
CP1 [07/31/14 15:53:36.225] 0x000e4609 00000000 GlobalRouter SW INFO <UNPRINTABLE>
rlogind: session 0 *IN USE* via user: \xc3\xb4\xb6\x8d\x0e\xb9+\xb1\xa2aL0~o\x3m\xf9\x8c
\x05g\xd2.\xc9pSWP\xc5\xd9 ip 10.139.82.29 <UNPRINTABLE>
```

The system encrypts AP information before writing it to the log file. The encrypted information is for debugging purposes. Only an Avaya Customer Service engineer can decrypt the information. ACLI commands display the logs without the encrypted information. Avaya recommends that you do not edit the log file.

The following table describes the system message severity levels.

Table 6: Severity levels

Severity level	Definition
INFO	Information only. Requires no action.
ERROR	A nonfatal condition occurred. You can be required to take appropriate action. For example, the system generates an error message if it is unable to lock onto the semaphore required to initialize the IP addresses used to transfer the log file to a remote host.
WARNING	A nonfatal condition occurred. No immediate action is needed.
FATAL	A fatal condition occurred. The system cannot recover without restarting. For example, a fatal message is generated after the configuration database is corrupted.

Based on the severity code in each message, the platform dispatches each message to one or more of the following destinations:

- workstation display
- local log file
- one or more remote hosts

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration except in the following case. When you configure the system log table, under the System Log Table tab, you must select either IPv4 or IPv6.

Internally, Virtual Services Platform 9000 has four severity levels for log messages: INFO, WARNING, ERROR, and FATAL. The system log supports eight different severity levels:

- Debug
- Info
- Notice
- Warning
- Critical
- Error
- Alert
- Emergency

The following table shows the default mapping of internal severity levels to syslog severity levels.

Table 7: Default and system log severity level mapping

UNIX system error codes	System log severity level	Internal VSP 9000 severity level
0	Emergency	Fatal
1	Alert	—
2	Critical	—
3	Error	Error
4	Warning	Warning
5	Notice	—
6	Info	Info
7	Debug	—

Log files

The log file captures hardware and software log messages, and alarm messages. Virtual Services Platform 9000 can log to the external flash. Avaya strongly recommends that you configure logging to an external flash and keep an external card in each CP module at all times. The system supports

2 GB Compact Flash cards. By default, the system logs to the external flash. If the external flash does not exist or the system configuration does not log to the external flash, the system logs to the internal flash instead.

To log to a file on the external or the internal flash, the used disk space on the flash must be below 75%. If the used disk space of the flash is more than 75%, the system stops logging to a file on the flash and raises an alarm even though the system always saves logs in the internal memory. The system saves internal log messages in a circular list in the memory, which overwrite older log messages as the log fills. Unlike the log messages in a log file, the internal log messages in the memory do not contain encrypted information, which can limit the information available during troubleshooting. Free up the disk space on the flash if the system generates the disk space 75% full alarm. After the disk space utilization returns below 75%, the system clears the alarm, and then starts logging to a file again.

Log file naming conventions

The following list provides the naming conventions for the log file:

- The log file format is log.xxxxxxx.sss. The prefix of the log file name is log. The six characters after the log file prefix contain the last three bytes of the chassis base MAC address. The next two characters specify the slot number of the CP module that generated the logs. The last three characters (sss) denote the sequence number of the log file.
- The system increments the sequence number of the log file for each new log file created after the existing log file reaches the maximum configured size.
- At initial system start up when no log file exists, the system creates a new log file with the sequence number 000. After a restart, the system finds the newest log file from both the external flash and the internal flash based on file timestamps. If the newest log file is on the flash that is used for logging, the system continues to use the newest log file for logging. If the newest log file exists on the flash that is not used for logging, the system creates a new log file with an incremented sequence number on the flash that is used for logging.

Log file transfer

The system logs contain important information for debugging and maintaining Virtual Services Platform 9000. After the current log file reaches the configured maximum size, the system creates a new log file for logging. The system transfers old log files to a remote host. You can configure up to 10 remote hosts, which creates long-term backup storage of your system log files.

Of the 10 configured remote hosts, 1 is the primary host and the other 9 are redundant. Upon initiating a transfer, system messaging attempts to use host 1 first. If host 1 is not reachable, system messaging tries hosts 2 to 10.

If the log file transfer is unsuccessful, the system keeps the old log files on the external flash or the internal flash. The system attempts to transfer old log files after the new log file reaches the configured maximum size. The system also attempts to transfer old log files periodically (once in one hundred log writes) if the disk space on the flash is more than 75% full.

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration.

You can specify the following information to configure the transfer criteria:

- The maximum size of the log file.

- The IP address of the remote host.
- The name prefix of the log file to store on the remote host.

The system appends a suffix of .xxxxxxx.sss to the file name. The first six characters of the suffix contain the last three bytes of the chassis base MAC address. The next two characters specify the slot number of the CP module that generated the logs. The last three characters (sss) denote the sequence number of the log file. For example, if you configure the name prefix as mylog, a possible file name is mylog.90000001.001.

- The user name and password, if using File Transfer Protocol (FTP) for file transfer. Use the following commands to configure the user name and password:

```
boot config host user WORD<0-16>
```

```
boot config host password WORD<0-16>
```

Be aware of the following restrictions to transfer log files to a remote host:

- The remote host IP address must be reachable.
- If you transfer a log file from a host to the system, (for example, to display it with a show command), rename the log file. Failure to rename the log file can cause the system to use the recently transferred file as the current log, if the sequence number in the extension is higher than the current log file. For example, if bf860005.002 is the current log file and you transfer bf860005.007 to the system, the system logs future messages to the bf860005.007 file. You can avoid this if you rename the log file to something other than the format used by system messaging.
- If your TFTP server is a UNIX-based machine, files written to the server must already exist. For example, you must create dummy files with the same names as your system logs. This action is commonly performed by using the touch command (for example, `touch bf860005.001`).

Three parameters exist to configure the log file:

- The minimum acceptable free space available on flash for logging.
- The maximum size of the log file.
- The percentage of free disk space the system can use for logging.

Although these three parameters exist, you can only configure the maximum size of the log file. Virtual Services Platform 9000 does not support the minimum size and percentage of free disk space parameters. The flash must be less than 75% full for the system to log a file. If the flash is more than 75% full, logging to a file stops to prevent exhausting disk space.

Log configuration using ACLI

Use log files and messages to perform diagnostic and fault management functions.

Configuring a UNIX system log and syslog host

Configure the syslog to control a facility in UNIX machines that logs SNMP messages and assigns each message a severity level based on importance.

About this task

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration using ACLI.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable the system log:

```
syslog enable
```

3. Specify the IP header in syslog packets:

```
syslog ip-header-type <circuitless-ip|default|management-virtual-ip>
```

4. Configure the maximum number of syslog hosts:

```
syslog max-hosts <1-10>
```

5. Create the syslog host:

```
syslog host <1-10>
```

6. Configure the IP address for the syslog host:

```
syslog host <1-10> address WORD <0-46>
```

7. Enable the syslog host:

```
syslog host <1-10> enable
```

Configure optional syslog host parameters by using the variables in the following variable definition tables.

8. View the configuration to ensure it is correct:

```
show syslog [host <1-10>]
```

Example

Configure a UNIX system log host address to IPv4 address 192.0.2.52 and syslog host:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:~(config)#syslog enable
VSP-9012:~(config)#syslog host 1 address 192.0.2.52
VSP-9012:~#syslog host 1 enable
VSP-9012:1(config)#show syslog host 1
```



```

        Id : 1
        IpAddr : 192.0.2.52
        UdpPort : 515
        Facility : local7
        Severity : info|warning|error|fatal
        MapInfoSeverity : info
        MapWarningSeverity : warning
        MapErrorSeverity : error
        MapMfgSeverity : notice
        MapFatalSeverity : emergency
        Enable : true

```

```
VSP-9012:1(config)#show syslog
```

```

Enable      : true
Max Hosts   : 5
OperState   : active
             header : default
Total number of configured hosts : 1
Total number of enabled hosts : 1
Configured host : 1
Enabled host : 1

```

Configure a UNIX system log host address to IPv6 address 2001:DB8:: and syslog host:

```

VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:#1(config)#syslog host 2 address 2001:DB8:: udp-port 515
VSP-9012:1(config)#syslog host 2 udp-port 515
VSP-9012:1(config)#syslog host 2 enable
VSP-9012:1(config)#show syslog host 2

```

```

        Id : 2
        IpAddr : 2001:DB8::
        UdpPort : 515
        Facility : local7
        Severity : info|warning|error|fatal
        MapInfoSeverity : info
        MapWarningSeverity : warning
        MapErrorSeverity : error
        MapMfgSeverity : notice
        MapFatalSeverity : emergency
        Enable : true

```

Variable definitions

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

Variable	Value
enable	Enables the sending of syslog messages on the device. The default is disabled. Use the no operator before this parameter, no syslog enable to disable the sending of syslog messages on the device. The default is enabled.
ip-header-type <circuitless-ip default management-virtual-ip>	Specifies the IP header in syslog packets to circuitless-ip, default, or management-virtual-ip. <ul style="list-style-type: none"> If the value is default, the IP address of the VLAN is used for syslog packets that transmit in-band using input/output (I/O) ports. For syslog packets that transmit out-of-band

Table continues...

Variable	Value
	<p>through the management port, the physical IP address of the master CPU is used in the IP header.</p> <ul style="list-style-type: none"> If the value is management-virtual-ip, the virtual management IP address of the device is used in the IP header for syslog packets that transmit out-of-band only through the management port. If the value is circuitless-ip, then for all syslog messages (in-band or out-of-band), the circuitless IP address is used in the IP header. If you configure multiple circuitless IPs, the first circuitless IP configured is used.
max-hosts <1-10>	Specifies the maximum number of syslog hosts supported, from 1–10. The default is 5.

Use the data in the following table to use the `syslog host` command.

Table 8: Variable definitions

Variable	Value
1–10	Creates and configures a host instance. Use the no operator before this parameter, no syslog host to delete a host instance.
address WORD <0–46>	Configures a host location for the syslog host. WORD <0–46> is the IPv4 or IPv6 address of the UNIX system syslog host in the format A.B.C.D or x:x:x:x:x:x. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration using ACLI.
enable	Enables the syslog host. Use the no operator before this parameter, no syslog host enable to disable syslog host. The default is disabled.
facility {local0 local1 local2 local3 local4 local5 local6 local7}	Specifies the UNIX facility in messages to the syslog host. {local0 local1 local2 local3 local4 local5 local6 local7} is the UNIX system syslog host facility. The default is local7.
maperror {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for error messages. The default is error.
mapfatal {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for fatal messages. The default is emergency.
mapinfo {emergency alert critical error warning notice info debug}	Specifies the syslog severity level to use for information messages. The default is info.
mapwarning {emergency alert critical error warning notice info debug}	Specifies the syslog severity to use for warning messages. The default is warning.
severity <info warning error fatal> [<info warning error fatal>] [<info warning error fatal>]	Specifies the severity levels for which to send syslog messages for the specified modules. The default is info.

Table continues...

Variable	Value
udp-port <514-530>	Specifies the User Datagram Protocol port number on which to send syslog messages to the syslog host. This value is the UNIX system syslog host port number from 514–530. The default is 514.

Configuring logging

Configure logging to determine the types of messages to log and where to store the messages.

About this task

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs ACLI log and SNMP log information regardless of the logging level you set. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Define which messages to log:


```
logging level <0-4>
```
3. Write the log file from memory to a file:


```
logging write WORD<1-1536>
```
4. Show logging on the screen:


```
logging screen
```

Example

Define which messages to log to 0 to record all messages. Write the log file from memory to file log2. Display logging on the screen.

```
VSP-9010:1>enable
VSP-9010:1#configure terminal
VSP-9010:1(config)#logging level 0
VSP-9010:1(config)#logging write log2
VSP-9010:1(config)#logging screen
```

Variable definitions

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

Variable	Value
level <0–4>	Shows and configures the logging level. The level is one of the following values: <ul style="list-style-type: none"> • 0: Information — Records all messages. • 1: Warning — Records only warning and more serious messages. • 2: Error — Records only error and more serious messages. • 3: Manufacturing — This parameter is not available for customer use. • 4: Fatal — Records only fatal messages.
logToExtFlash	Starts logging system messages to the external flash. The default logging location is the external flash device. Avaya recommends that you use logging to the external flash. Use the no form of the command to stop logging to external flash and log to internal flash instead: <code>no logging logToExtFlash</code>
screen	Configures the log display on the screen to on. Use the no form of the command to stop the log display on the screen: <code>no logging screen</code>
transferFile <1–10> address {A.B.C.D} filename-prefix WORD<0–200	Transfers the syslog file to a remote FTP/TFTP server. <1–10> specifies the file ID. The address {A.B.C.D} option specifies the IP address. The filename-prefix WORD<0–200> option sets the filename prefix for the log file at the remote host.
write WORD<1-1536>	Writes the log file with the designated string. WORD<1-1536> is the string or command that you append to the log file. If the string contains spaces, you must enclose the string in quotation marks (").

Configuring the remote host address for log transfer

Configure the remote host address for log transfer. The system transfers the current log file to a remote host after the log file size reaches the maximum size.

Before you begin

- The IP address you configure for the remote host must be reachable at the time of configuration.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```
2. Configure the remote host address for log transfer:

```
logging transferFile {1-10} address {A.B.C.D} [filename WORD<0-255>]
```

Example

Configure the remote host address for log transfer to 192.0.2.10:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#logging transferFile 1 address 192.0.2.10
```

Variable definitions

Use the data in the following table to use the **logging transferFile** command.

Variable	Value
1-10	Specifies the file ID to transfer.
address {A.B.C.D}	Specifies the IP address of the host to which to transfer the log file. The remote host must be reachable or the configuration fails.
filename-prefix WORD<0-255>	Specifies the name of the file on the remote host. If you do not configure a name, the current log file name is the default.

Configuring system logging to external storage

System logs are a valuable diagnostic tool. You can send log messages to external flash for later retrieval.

Before you begin

- You must install a CF card in the CP module before you can log to external storage.

Caution:

Risk of data loss

Before you remove the CF card from the master CP module, you must stop the logging of system messages. Failure to do so can corrupt the file system on the CF card and cause the log file to be permanently lost.

About this task

Define the maximum log file sizes to bound the file storage size on the Compact Flash (CF) card. The system transfers the current log file to a remote host after the log file size reaches the maximum size.

You can change log file parameters at any time without restarting the system. Changes made to these parameters take effect immediately.

Avaya recommends that you configure logging to an external flash and keep an external flash in each CP module at all times. If the external flash does not exist, the system raises an alarm, and then logs to the internal flash instead.

Procedure

1. Enter Global Configuration mode:

Log and trap information

```
enable
```

```
configure terminal
```

2. Enable system logging to a CF card:

```
boot config flags logging
```

3. Configure the logfile parameters:

```
boot config logfile <64-500> <500-16384> <10-90>
```

Example

Enable system logging to a CF card and configure the logfile parameters:

```
VSP-9010:1>enable
VSP-9010:1#configure terminal
VSP-9010:1(config)#boot config flags logging
VSP-9010:1(config)#boot config logfile 64 600 10
```

Variable definitions

Use the data in the following table to use the `boot config` command.

Variable	Value
flags logging	Enables or disables logging to a file on the external flash. The log file uses the naming format log.xxxxxxx.sss. The first six characters after the prefix of the file name log contain the last three bytes of the chassis base MAC address. The next two characters specify the slot number of the CP module that generated the logs. The last three characters denote the sequence number of the log file. The system generates multiple sequence numbers for the same chassis and same slot, if you replace or reinsert the CP module, or if the maximum log file size is reached.
logfile <64-500> <500-16384> <10-90>	Configures the logfile parameters <ul style="list-style-type: none">• <64-500> specifies the minimum free memory space on the external storage device from 64–500 KB. Virtual Services Platform 9000 does not support this parameter.• <500-16384> specifies the maximum size of the log file from 500–16384 KB.• <10-90> specifies the maximum percentage, from 10–90%, of space on the external storage device the logfile can use. Virtual Services Platform 9000 does not support this parameter.

Configuring system message control

Configure system message control to suppress duplicate error messages on the console, and to determine the action to take if they occur.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure system message control action:

```
sys msg-control action <both|send-trap|suppress-msg>
```

3. Configure the maximum number of messages:

```
sys msg-control max-msg-num <2-500>
```

4. Configure the interval:

```
sys msg-control control-interval <1-30>
```

5. Enable message control:

```
sys msg-control
```

Example

Configure system message control action to suppress the message. Configure the maximum number of messages to 10. Configure the message control interval to 15. Enable message control.

```
VSP-9010:1>enable
VSP-9010:1#configure terminal
VSP-9010:1(config)#sys msg-control action suppress-msg
VSP-9010:1(config)#sys msg-control max-msg-num 10
VSP-9010:1(config)#sys msg-control control-interval 15
VSP-9010:1(config)#sys msg-control
```

Variable definitions

Use the data in the following table to use the `sys msg-control` command.

Variable	Value
action <both send-trap suppress-msg>	Configures the message control action. You can either suppress the message or send a trap notification, or both. The default is suppress.
control-interval <1-30>	Configures the message control interval in minutes. The valid options are 1–30. The default is 5.
max-msg-num <2-500>	Configures the number of occurrences of a message after which the control action occurs. To configure the maximum number of occurrences, enter a value from 2–500. The default is 5.

Extending system message control

Use the force message control option to extend the message control feature functionality to the software and hardware log messages.

About this task

To enable the message control feature, you must specify an action, control interval, and maximum message number. After you enable the feature, the log messages, which get repeated and cross the maximum message number in the control interval, trigger the force message feature. You can either suppress the message or send a trap notification, or both.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the force message control option:

```
sys force-msg WORD<4-4>
```

Example

Configure the force message option. Add a force message control pattern. If you use a wildcard pattern (****), all messages undergo message control.

```
VSP-9010:1>enable
VSP-9010:1#configure terminal
VSP-9010:1(config)#sys force-msg ****
```

Variable definitions

Use the data in the following table to use the `sys force-msg` command.

Variable	Value
WORD<4-4>	Adds a forced message control pattern, where WORD<4-4> is a string of 4 characters. You can add a four-byte pattern into the force-msg table. The software and the hardware log messages that use the first four bytes that match one of the patterns in the force-msg table undergo the configured message control action. You can specify up to 32 different patterns in the force-msg table, including a wildcard pattern (****) as well. If you specify the wildcard pattern, all messages undergo message control.

Viewing logs

View log files by file name, category, severity, or CP module to identify possible problems. View ACLI command and SNMP trap logs, which the system logs as normal log messages and log to the system log file.

About this task

Unprintable characters in log and trace messages are encoded in C language hexadecimal notation in the string and surrounded with the <UNPRINTABLE> . . .<UNPRINTABLE> tag. For example, the device converts a 4-byte sequence of 0x01, 0x02, 0xfe, 0xff to \x01\x02\xfe\xff, and adds the

<UNPRINTABLE>...<UNPRINTABLE> tag to the log or trace message. The following example displays how the message appears:

```
CP1 [07/31/14 15:53:36.225] 0x000e4609 00000000 GlobalRouter SW INFO <UNPRINTABLE>
rlogind: session 0 *IN USE* via user: \xc3\xb4\xb6\x8d\x0e\xb9+\xb1\xa2aL0~o\xfb3m\xfb9\x8c
\x05g\xd2.\xc9pSWP\xc5\xd9 ip 10.139.82.29 <UNPRINTABLE>
```

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display log information:

```
show logging file [alarm][CPU WORD<0-100>] [event-code WORD<0-10>]
[module WORD<0-100>] [name-of-file WORD<1-99>] [save-to-file
WORD<1-99>] [severity WORD<0-25>] [tail] [vrf WORD<0-32>]
```

Example

Display log information:

```
VSP-9012:1>show logging file module cliilog
CP1 [08/21/11 14:29:57.231] 0x002c0600 00000000 GlobalRouter CLILOG INFO 1 CONSOLE
rwa en
CP1 [08/21/11 14:29:58.771] 0x002c0600 00000000 GlobalRouter CLILOG INFO 2 CONSOLE
rwa config t
CP1 [08/21/11 14:30:06.743] 0x002c0600 00000000 GlobalRouter CLILOG INFO 3 CONSOLE
rwa source basic.cfg
CP1 [08/21/11 14:30:07.018] 0x002c0600 00000000 GlobalRouter CLILOG INFO 4 CONSOLE
rwa config terminal
CP1 [08/21/11 14:30:07.026] 0x002c0600 00000000 GlobalRouter CLILOG INFO 5 CONSOLE
rwa boot config flags fabric-profile 1
CP1 [08/21/11 14:30:07.027] 0x002c0600 00000000 GlobalRouter CLILOG INFO 6 CONSOLE
rwa boot config flags ftpd
CP1 [08/21/11 14:30:07.028] 0x002c0600 00000000 GlobalRouter CLILOG INFO 7 CONSOLE
rwa boot config flags rlogind
CP1 [08/21/11 14:30:07.029] 0x002c0600 00000000 GlobalRouter CLILOG INFO 8 CONSOLE
rwa boot config flags sshd
CP1 [08/21/11 14:30:07.030] 0x002c0600 00000000 GlobalRouter CLILOG INFO 9 CONSOLE
rwa boot config flags telnetd
CP1 [08/21/11 14:30:07.031] 0x002c0600 00000000 GlobalRouter CLILOG INFO 10 CONSOLE
rwa cli timeout 65535
CP1 [08/21/11 14:30:07.032] 0x002c0600 00000000 GlobalRouter CLILOG INFO 11 CONSOLE
rwa password password-history 3
CP1 [08/21/11 14:30:07.033] 0x002c0600 00000000 GlobalRouter CLILOG INFO 12 CONSOLE
rwa cliilog enable
CP1 [08/21/11 14:30:07.034] 0x002c0600 00000000 GlobalRouter CLILOG INFO 13 CONSOLE
rwa snmplog enable
CP1 [08/21/11 14:30:07.036] 0x002c0600 00000000 GlobalRouter CLILOG INFO 14 CONSOLE
rwa no sys ecn-compatibility
CP1 [08/21/11 14:30:07.046] 0x002c0600 00000000 GlobalRouter CLILOG INFO 15 CONSOLE
rwa interface mgmtEthernet 1/1
CP1 [08/21/11 14:30:07.047] 0x002c0600 00000000 GlobalRouter CLILOG INFO 16 CONSOLE
rwa ip address 47.17.159.49 255.255.255.0
CP1 [08/21/11 14:30:07.049] 0x002c0600 00000000 GlobalRouter CLILOG INFO 17 CONSOLE
rwa exit
CP1 [08/21/11 14:30:07.050] 0x002c0600 00000000 GlobalRouter CLILOG INFO 18 CONSOLE
rwa interface GigabitEthernet 10/11
CP1 [08/21/11 14:30:07.051] 0x002c0600 00000000 GlobalRouter CLILOG INFO 19 CONSOLE
rwa no shutdown
CP1 [08/21/11 14:30:07.053] 0x002c0600 00000000 GlobalRouter CLILOG INFO 20 CONSOLE
rwa exit
CP1 [08/21/11 14:30:07.054] 0x002c0600 00000000 GlobalRouter CLILOG INFO 21 CONSOLE
rwa interface gigabitethernet 10/11
```

Log and trap information

```
CP1 [08/21/11 14:30:07.056] 0x002c0600 00000000 GlobalRouter CLILOG INFO 22 CONSOLE
rwa ipv6 interface vlan 3
CP1 [08/21/11 14:30:07.079] 0x002c0600 00000000 GlobalRouter CLILOG INFO 23 CONSOLE
rwa ipv6 interface enable
```

Variable definitions

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

Variable	Value
alarm	Displays alarm log entries.
CPU <i>WORD</i> <0-100>	Filters and lists the logs according to the CP module that generated the message. Specify a string length of 0–100 characters. To specify multiple filters, separate each CP module by the vertical bar (), for example, <code>show logging file CPU CP1 CP2 IO1</code> . Following are some of the available CPU qualifiers: <ul style="list-style-type: none">• CP1• CP2• IO1• IO2• SF1• SF6
event-code <i>WORD</i> <0–10>	Specifies a number that precisely identifies the event reported.
module <i>WORD</i> <0-100>	Filters and lists the logs according to module. Specifies a string length of 0–100 characters. Categories include SNMP, EAP, RADIUS, RMON, WEB, IGMP, HW, MLT, FILTER, QOS, CLILOG, SW, CPU, IP, VLAN, IPMC, IP-RIP, OSPF, PIM, POLICY, RIP and SNMPLOG. To specify multiple filters, separate each category by the vertical bar (), for example, <code>OSPF FILTER QOS</code> .
name-of-file <i>WORD</i> <1-99>	Displays the valid logs from this file. For example, <code>/intflash/logcopy.txt</code> . You cannot use this command on the current log file—the file into which the messages currently log. Specify a string length of 1–99 characters.
save-to-file <i>WORD</i> <1-99>	Redirects the output to the specified file and removes all encrypted information. You cannot use the tail option with the save-to-file option. Specify a string length of 1–99 characters. The format for the file name is: <code>/intflash/<filename></code> , <code>/extflash/<filename></code> , or <code>/usb/<filename></code> .
severity <i>WORD</i> <0-25>	Filters and lists the logs according to severity. Choices include INFO, ERROR, WARNING, and FATAL. To specify multiple filters, separate each severity by the vertical bar (), for example, <code>ERROR WARNING FATAL</code> .
tail	Shows the last results first.
vrf <i>WORD</i> <0–32>	Specifies the name of a VRF instance to show log messages that only pertain to that VRF.

Configuring ACLI logging

Use ACLI logging to track all ACLI commands executed for archiving and fault management purposes. You can track system changes made in the ACLI. The system logs ACLI commands to the system log file as CLILOG module.

* Note:

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs ACLI log and SNMP log information regardless of the logging level you set. This is not the case for other INFO messages.

About this task

The log captures the ACLI command information in the following format:

```
CP1 [08/21/11 14:30:07.028] 0x002c0600 00000000 GlobalRouter CLILOG INFO
7 CONSOLE rwa boot config flags rlogind
```

The following list identifies the relevant ACLI command information in the preceding log message:

- [08/21/11 14:30:07.028] — The command execution time.
- CONSOLE — The source of the connection. If the connection is a Telnet connection, the message also provides the IP address associated with the connection.
- rwa — The user ID that used the command.
- boot config flags rlogind — The command used.

Unprintable characters in log and trace messages are encoded in C language hexadecimal notation in the string and surrounded with the <UNPRINTABLE>...<UNPRINTABLE> tag. For example, the device converts a 4-byte sequence of 0x01, 0x02, 0xfe, 0xff to \x01\x02\xfe\xff, and adds the <UNPRINTABLE>...<UNPRINTABLE> tag to the log or trace message. The following example displays how the message appears:

```
CP1 [07/31/14 15:53:36.225] 0x000e4609 00000000 GlobalRouter SW INFO <UNPRINTABLE>
rlogind: session 0 *IN USE* via user: \xc3\xb4\xb6\x8d\x0e\xb9+\xb1\xa2aLO~o\x3m\xf9\x8c
\x05g\xd2.\xc9pSWP\xc5\xd9 ip 10.139.82.29 <UNPRINTABLE>
```

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Enable ACLI logging:


```
cliolog enable
```
3. Disable ACLI logging:


```
no cliolog enable
```
4. Ensure that the configuration is correct:

```
show clilog
```

5. View the ACLI log:

a. View log files generated by Release 3.2 and greater:

```
show logging file module clilog
```

b. View log files generated by releases prior to Release 3.2:

```
show clilog file [grep WORD<1-256>] [tail]
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#clilog enable
VSP-9012:1(config)#show logging file module clilog
CP1 [08/21/11 14:29:57.231] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 1 CONSOLE rwa en
CP1 [08/21/11 14:29:58.771] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 2 CONSOLE rwa config t
CP1 [08/21/11 14:30:06.743] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 3 CONSOLE rwa source basic.cfg
CP1 [08/21/11 14:30:07.018] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 4 CONSOLE rwa config terminal
CP1 [08/21/11 14:30:07.026] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 5 CONSOLE rwa boot config flags fabric-profile 1
CP1 [08/21/11 14:30:07.027] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 6 CONSOLE rwa boot config flags ftpd
CP1 [08/21/11 14:30:07.028] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 7 CONSOLE rwa boot config flags rlogind
CP1 [08/21/11 14:30:07.029] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 8 CONSOLE rwa boot config flags sshd
CP1 [08/21/11 14:30:07.030] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 9 CONSOLE rwa boot config flags telnetd
CP1 [08/21/11 14:30:07.031] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 10 CONSOLE rwa cli timeout 65535
CP1 [08/21/11 14:30:07.032] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 11 CONSOLE rwa password password-history 3
CP1 [08/21/11 14:30:07.033] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 12 CONSOLE rwa clilog enable
CP1 [08/21/11 14:30:07.034] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 13 CONSOLE rwa snmplog enable
CP1 [08/21/11 14:30:07.046] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 15 CONSOLE rwa interface mgmtEthernet 1/1
CP1 [08/21/11 14:30:07.047] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 16 CONSOLE rwa ip address 192.0.2.49 255.255.255.0
CP1 [08/21/11 14:30:07.049] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 17 CONSOLE rwa exit
CP1 [08/21/11 14:30:07.050] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 18 CONSOLE rwa interface GigabitEthernet 10/11
CP1 [08/21/11 14:30:07.051] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 19 CONSOLE rwa no shutdown
CP1 [08/21/11 14:30:07.053] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 20 CONSOLE rwa exit
CP1 [08/21/11 14:30:07.054] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 21 CONSOLE rwa interface gigabitethernet 10/11
CP1 [08/21/11 14:30:07.056] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 22 CONSOLE rwa ipv6 interface vlan 3
CP1 [08/21/11 14:30:07.079] 0x002c0600 00000000 GlobalRouter CLILOG
INFO 23 CONSOLE rwa ipv6 interface enable
```

Variable definitions

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

Variable	Value
enable	Activates ACLI logging. To disable, use the <code>no cliilog enable</code> command.

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

*** Note:**

The `show cliilog file` command only applies to log files generated by releases prior to Release 3.2.

Variable	Value
tail	Shows the last results first.
grep WORD<1-256>	Performs a string search in the log file. <i>WORD<1-256></i> is the string, of up to 256 characters in length, to match.

Log configuration using EDM

Use log files and messages to perform diagnostic and fault management functions. This section provides procedures to configure and use the logging system in Enterprise Device Manager (EDM).

Configuring the system log

Configure the system log to track all user activity on the device. The system log can send messages to up to ten syslog hosts.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **System Log**.
3. In the **System Log** tab, select **Enable**.
4. Configure the maximum number of syslog hosts.
5. Configure the IP header type for the syslog packet.
6. Click **Apply**.

System Log field descriptions

Use the data in the following table to use the **System Log** tab.

Name	Description
Enable	Enables or disables the syslog feature. If you select this variable, this feature sends a message to a server on a network that is configured to receive and store diagnostic messages from this device. You can configure the type of messages sent. The default is enabled.
MaxHosts	Specifies the maximum number of remote hosts considered active and can receive messages from the syslog service. The range is 0–10 and the default is 5.
OperState	Specifies the operational state of the syslog service. The default is active.
Header	<p>Specifies the IP header in syslog packets to circuitlessIP, default, or managementVIP.</p> <ul style="list-style-type: none"> • If the value is default, the IP address of the system uses the VLAN for syslog packets that transmit in-band using input/output (I/O) ports. For syslog packets that transmit out-of-band through the management port, the system uses the physical IP address of the master CPU in the IP header. • If the value is managementVIP, the system uses the virtual management IP address of the device in the IP header for syslog packets that transmit out-of-band only through the management port. • If the value is circuitlessIP, the system uses the circuitless IP address in the IP header for all syslog messages (in-band or out-of-band). If you configure multiple circuitless IPs, the system uses the first circuitless IP that you configure. <p>The default value is default.</p>

Configuring the system log table

Use the system log table to customize the mappings between the severity levels and the type of alarms, and to configure an entry for a remote syslog server. Virtual Services Platform 9000 supports up to 10 syslog servers.

About this task

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration except in the following case. When you configure the system log table, under the **System Log Table** tab, you must select **ipv4** or **ipv6**, in the **AddressType** box. The **Address** box supports both IPv4 and IPv6 addresses.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **System Log**.

3. Click the **System Log Table** tab.
4. Click **Insert**.
5. Configure the parameters as required.
6. Click **Insert**.
7. To modify mappings, double-click a parameter to view a list of options.
8. Click **Apply**.

System Log Table field descriptions

Use the data in the following table to use the **System Log Table** tab.

Name	Description
Id	Specifies the ID for the syslog host. The range is 1–10.
AddressType	Specifies if the address is an IPv4 or an IPv6 address.
Address	Specifies the IP address of the syslog host. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses.
UdpPort	Specifies the UDP port to use to send messages to the syslog host (514–530). The default is 514.
Enable	Enables or disables the sending of messages to the syslog host. The default is disabled.
HostFacility	Specifies the syslog host facility used to identify messages (LOCAL0 to LOCAL7). The default is LOCAL7.
Severity	Specifies the message severity for which the system sends syslog messages. The default is INFO.
MapInfoSeverity	Specifies the syslog severity to use for INFO messages. The default is INFO.
MapWarningSeverity	Specifies the syslog severity to use for WARNING messages. The default is WARNING.
MapErrorSeverity	Specifies the syslog severity to use for ERROR messages. The default is ERROR.
MapFatalSeverity	Specifies the syslog severity to use for FATAL messages. The default is EMERGENCY.
MapMfgSeverity	Specifies the syslog severity to use for Accelar manufacturing messages. The default is ERROR.

SNMP trap configuration using ACLI

Use Simple Network Management Protocol (SNMP) traps and notifications to gather information about device activities, alarms, and other information on management stations.

For more information about how to configure SNMP community strings and related topics, see *Avaya Virtual Services Platform 9000 Security*, NN46250–601.

Configuring an SNMP host

Configure an SNMP host to enable the system can forward SNMP traps to a host for monitoring. You can use SNMPv1, SNMPv2c, or SNMPv3.

About this task

You configure the target table parameters (security name and model) as part of the host configuration.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure an SNMPv1 host:

```
snmp-server host WORD<1-256> [port <1-65535>] v1 WORD<1-32> [filter WORD<1-32>]
```

3. Configure an SNMPv2c host:

```
snmp-server host WORD<1-256> [port <1-65535>] v2c WORD<1-32> [inform [timeout <0-2147483647>] [retries <0-255>] [mms <0-2147483647>]] [filter WORD<1-32>]
```

4. Configure an SNMPv3 host:

```
snmp-server host WORD<1-256> [port <1-65535>] v3 {noAuthNoPriv|authNoPriv|AuthPriv} WORD<1-32> [inform [timeout <0-2147483647>] [retries <0-255>] [filter WORD<1-32>]
```

5. Ensure that the configuration is correct:

```
show snmp-server host
```

Example

1. Configure the target table entry:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#snmp-server host 198.202.188.207 port 162 v2c ReadView inform
timeout 1500 retries 3 mms 484
```

2. Configure an SNMPv3 host:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#snmp-server host 4717:0:0:0:0:0:7933:6 port 163 v3 authPriv
Lab3 inform timeout 1500 retries 3
```


Variable definitions

Use the data in the following table to use the `snmp-server host` command.

Variable	Value
inform [timeout <0-2147483647>] [retries <0-255>] [mms <0-2147483647>]	Sends SNMP notifications as inform (rather than trap). To use all three options in one command, you must use them in the following order: <ol style="list-style-type: none"> 1. timeout <0-2147483647>—Specifies the time to wait for a reply before resending the inform message. Time is specified in centiseconds. 2. retries <0-255>—Specifies the number of packets to be sent if no reply is received 3. mms <0-2147483647>—Specifies the maximum message size as an integer with a range of 0–2147483647.
filter WORD<1-32>	Specifies the filter profile to use.
noAuthNoPriv authNoPriv AuthPriv	Specifies the security level.
port <1-65535>	Specifies the port number that will be set as the destination port at the UDP level in the trap packet.
WORD<1-32>	Specifies the security name, which identifies the principal that generates SNMP messages.
WORD<1-256>	Specifies either an IPv4 or IPv6 address. <p>* Note:</p> <p>The SNMP server host IPv6 format should be x:x:x:x:x:x. Avaya recommends you do not use :: in the IPv6 address. If you use :: the port number becomes part of the IPv6 address in the SNMP target address table.</p>

Configuring an SNMP notify filter table

Configure the notify table to select management targets to receive notifications, as well as the type of notification to send to each management target.

Before you begin

- For more information about the notify filter table, see RFC3413.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a new notify filter table:

```
snmp-server notify-filter WORD<1-32> WORD<1-32>
```

3. Ensure that the configuration is correct:

```
show snmp-server notify-filter
```

Example

Create a new notify filter table:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#snmp-server notify-filter profile3 99.3.4.1.4.3.1.1.4.1
VSP-9012:1#show snmp-server notify-filter
```

```
=====
Notify Filter Configuration
=====
Profile Name          Subtree              Mask
-----
profile1              +99.3.4.1.4.3.1.1.4.1  0x7f
profile2              +99.3.4.1.4.3.1.1.4.1  0x7f
profile3              +99.3.4.1.4.3.1.1.4.1  0x7
```

Variable definitions

Use the data in the following table to use the `snmp-server notify-filter` command.

Variable	Value
<code>WORD<1-32> WORD<1-32></code>	<p>Creates a notify filter table.</p> <p>The first instance of <code>WORD<1-32></code> specifies the name of the filter profile with a string length of 1–32.</p> <p>The second instance of <code>WORD<1-32></code> identifies the filter subtree OID with a string length of 1–32.</p> <p>If the subtree OID parameter uses a plus sign (+) prefix (or no prefix), this indicates include. If the subtree OID uses the minus sign (–) prefix, it indicates exclude.</p> <p>You do not calculate the mask because it is automatically calculated. You can use the wildcard character, the asterisk (*), to specify the mask within the OID. You do not need to specify the OID in the dotted decimal format; you can alternatively specify that the MIB parameter names and the OIDs are automatically calculated.</p>

Configuring SNMP interfaces

Configure an interface to send SNMP traps.

About this task

If Avaya Virtual Services Platform 9000 has multiple interfaces, configure the IP interface from which the SNMP traps originate.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the destination and source IP addresses for SNMP traps:

```
snmp-server sender-ip {A.B.C.D} {A.B.C.D}
```

3. If required, send the source address (sender IP) as the sender network in the notification message:

```
snmp-server force-trap-sender enable
```

4. If required, force the SNMP and IP sender flag to use the same value:

```
snmp-server force-iphdr-sender enable
```

5. Activate the generation of authentication traps:

```
snmp-server authentication-trap enable
```

Example

Configure the destination address to 192.0.2.2 and the source address to 192.0.2.5 and enable the generation of authentication traps:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#snmp-server sender-ip 192.0.2.2 192.0.2.5
VSP-9012:1(config)#snmp-server authentication-trap enable
```

Variable definitions

Use the data in the following table to use the **snmp-server** command.

Variable	Value
authentication-trap enable	Activates the generation of authentication traps.
community <i>WORD</i> <1–32>	Specifies a community string to create a community to use in forming a relationship between an SNMP agent and one or more SNMP managers. You require SNMP community strings to access the system using an SNMP-based management software. Use the no option to delete the community string: no snmp-server community <i>WORD</i> <1–32>
contact <i>WORD</i> <0–255>	Changes the sysContact information for Virtual Services Platform 9000. <i>WORD</i> <0–255> is an ASCII string from 0–255 characters (for example a phone extension or e-mail address).
force-iphdr-sender enable	Automatically configures the SNMP and IP sender to the same value. The default is disabled.

Table continues...

Variable	Value
force-trap-sender enable	Sends the configured source address (sender IP) as the sender network in the notification message.
group <i>WORD</i> <1–32> [<i>WORD</i> <0–32>] [auth-no-priv auth-priv no-auth-no-priv] [notify-view <i>WORD</i> <0–32> read-view <i>WORD</i> <0–32> write-view <i>WORD</i> <0–32>]	Creates a new user group. <ul style="list-style-type: none"> auth-no-priv auth-priv no-auth-no-priv – Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the no-auth-no-priv parameter is included, it creates three entries, one for SNMPv1 access, one for SNMPv2c access, and one for SNMPv3 access. <i>WORD</i><1–32> <i>WORD</i><1–32> – The first <i>WORD</i><1–32> specifies the group name for data access. The second <i>WORD</i><1–32> specifies the context name. notify-view <i>WORD</i><0–32> read-view <i>WORD</i><0–32> write-view <i>WORD</i><0–32> – Specifies the view name. The switch uses elements from the notify-view in the trap messages sent for the communities associated with this group. The switch makes elements from the read-view available for reading for the communities associated with this group. The switch makes elements from the write-view available to be modified by the communities associated with this group.
host <i>WORD</i> <1–256>[port<1–65535>][v1 v2c v3][<i>WORD</i> <1–32>][filter <i>WORD</i> <1–32>][inform][mms<1-2147483647>][retries<0-255>][timeout<1-2147483647>][noAuthPriv authNoPriv authPriv]	Configures hosts to receive SNMP notifications. <ul style="list-style-type: none"> host <i>WORD</i><1-256>— Specifies the IPv4 or IPv6 host address. port <1-65535>—Specifies the port number. v1 <i>WORD</i><1-32>—Specifies the SNMP v1 security name. v2c <i>WORD</i><1-32>—Specifies the SNMPv2 security name. inform—Specifies the notify type. timeout <1-2147483647>—Specifies the timeout value. retries <0-255>—Specifies the number of retries. mms <1-2147483647>—Specifies the maximum message size. v3 —Specifies SNMPv3. noAuthPriv authNoPriv authPriv —Specifies the security level. <i>WORD</i><1-32>—Specifies the user name. filter—Specifies a filter profile name.
location <i>WORD</i> <0–255>	Configures the sysLocation information for the system. < <i>WORD</i> 0-255> is an ASCII string from 0–255 characters.
login-success-trap enable	Enables the generation of login-success traps.

Table continues...

Variable	Value
name <i>WORD</i> <0–255>	Configures the sysName information for the system. < <i>WORD</i> 0-255> is an ASCII string from 0–255 characters.
notify-filter <i>WORD</i> <1–32>	Creates a new entry in the notify filter table. The first <i>WORD</i> <1-32> specifies the filter profile name, and the second <i>WORD</i> <1-32> specifies the subtree OID.
sender-ip <A.B.C.D> <A.B.C.D>	Configures the SNMP trap receiver and source IP addresses. Specify the IP address of the destination SNMP server that receives the SNMP trap notification in the first IP address. Specify the source IP address as the second IP address. Specify the source IP address of the SNMP trap notification packet that is transmitted in the second IP address. If this address is 0.0.0.0, the system uses the IP address of the local interface that is closest (from an IP routing table perspective) to the destination SNMP server.
user [engine-id <i>WORD</i> <16–97>] [group <i>WORD</i> <1–32>][notify-view <i>WORD</i> <0–32> read-view <i>WORD</i> <0–32> write-view <i>WORD</i> <0–32>] [<i>WORD</i> <1–32>] [md5 sha <i>WORD</i> <1–32>] [aes des <i>WORD</i> <1–32>]	Creates a new user in the USM table to authorize a user on a particular SNMP engine. <ul style="list-style-type: none"> • [aes des] <i>WORD</i><1–32> – Specifies a privacy protocol. If no value is entered, no authentication capability exists. The choices are aes and des. <i>WORD</i><1-32> assigns a privacy password. If no value is entered, no privacy capability exists. The range is 1–32 characters. You must set authentication before you can set the privacy option. • engine-id <i>WORD</i><1–32> – Assigns an SNMPv3 engine ID. The range is 10–64 characters. Use the no operator to remove this configuration. • group <i>WORD</i><1–32> – Specifies the group access name. • [md5 sha] <i>WORD</i><1–32> – Specifies an authentication protocol. If no value is entered, no authentication capability exists. The protocol choices are: MD5 and SHA. <i>WORD</i><1-32> specifies an authentication password. If no value is entered, no authentication capability exists. The range is 1–32 characters. • notify-view <i>WORD</i><0–32> read-view <i>WORD</i><0–32> write-view <i>WORD</i><0–32> – Specifies the view name.
view <i>WORD</i> <1–32>[<i>WORD</i> <1–32>]	<i>WORD</i> <1–32> specifies a new entry with this group name. The range is 1-32 characters. <i>WORD</i> <1–32> <i>WORD</i> <1–32> specifies the prefix that defines the set of MIB objects accessible by this SNMP entity. The range is 1-32 characters.

Enabling SNMP trap logging

Use SNMP trap logging to send a copy of all traps to the syslog server.

Before you begin

- You must configure and enable the syslog server.

About this task

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs ACLI log and SNMP log information regardless of the logging level you set. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable SNMP trap logging:

```
snmplog enable
```

3. Disable SNMP trap logging:

```
no snmplog enable
```

4. View the contents of the SNMP log:

- a. View the SNMP log files generated for Release 3.2 and greater:

```
show logging file module snmplog
```

- b. View the SNMP log files generated by releases prior to Release 3.2:

```
show snmplog [file [grep WORD<1-255>|tail]]
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSP-9012:1(config)#snmplog enable
VSP-9012:1(config)#show logging file module snmplog
CP1 [04/01/13 17:16:04.130] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 1
  ver=v2c private rcSysActionL1.0 = 7
CP1 [04/02/13 10:50:41.122] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 2
  ver=v2c public rcVrfRpTrigger.3 = 4
CP1 [04/02/13 14:22:11.620] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 3
  ver=v2c private rcSysActionL1.0 = 7
CP1 [04/03/13 11:03:35.991] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 4
  ver=v2c public pingCtlRowStatus.11.111.119.110.101.114.105.110.100.101.120.4
  9.8.116.101.115.116.105.112.118.52 = 4
CP1 [04/03/13 11:03:35.992] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 5
  ver=v2c public pingCtlTargetAddressType.11.111.119.110.101.114.105.110.100.1
  01.120.49.8.116.101.115.116.105.112.118.52 = 1
CP1 [04/03/13 11:03:35.992] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 6
  ver=v2c public pingCtlTargetAddress.11.111.119.110.101.114.105.110.100.101.1
  20.49.8.116.101.115.116.105.112.118.52 = 1.1.1.1
CP1 [04/03/13 11:03:35.993] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 7
  ver=v2c public pingCtlAdminStatus.11.111.119.110.101.114.105.110.100.101.120
  .49.8.116.101.115.116.105.112.118.52 = 2
CP1 [04/03/13 11:03:35.993] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 8
  ver=v2c public pingCtlDataSize.11.111.119.110.101.114.105.110.100.101.120.49
```

```
.8.116.101.115.116.105.112.118.52 = 16
--More-- (q = quit)
```

Variable definitions

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

Variable	Value
enable	Enables the logging of traps. Use the command <code>no snmplog enable</code> to disable the logging of traps.
file [grep WORD<1–255> tail]	The parameter only applies to log files generated by releases prior to Release 3.2: Shows the trap log file stored on external flash. You can optionally specify search or display parameters: <ul style="list-style-type: none"> • <code>grep WORD<1–255></code> performs a string search in the log file. <code>WORD<1–255></code> is the string, of up to 255 characters in length, to match. • <code>tail</code> shows the last results first.

SNMP trap configuration using EDM

Use Simple Network Management Protocol (SNMP) traps and notifications to gather information about device activities, alarms, and other information on management stations. This section provides procedures to configure and use SNMP traps with Enterprise Device Manager (EDM).

For information about how to configure SNMP community strings and related topics, see *Avaya Virtual Services Platform 9000 Security*, NN46250–601.

Configuring an SNMP host target address

Configure a target table to specify the list of transport addresses to use in the generation of SNMP messages.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **Target Table**.
3. In the **Target Table** tab, click **Insert**.
4. In the **Name** box, type a unique identifier.

5. In the **TDomain** box, select the transport type of the address. Select either **ipv4Tdomain** or **ipv6Tdomain**.
6. In the **TAddress** box, type the transport address and User Datagram Protocol (UDP) port.
7. In the **Timeout** box, type the maximum round trip time.
8. In the **RetryCount** box, type the number of retries to be attempted.
9. In the **TagList** box, type the list of tag values.
10. In the **Params** box, click the ellipsis (...) to select **TparamV1** or **TparamV2**.
11. Click **OK**.
12. In the **TMask** box, type the mask.
13. In the **MMS** box, type the maximum message size.
14. Click **Insert**.

Target Table field descriptions

Use the data in the following table to use the **Target Table** tab.

Name	Description
Name	Specifies a unique identifier for this table. The name is a community string.
TDomain	Specifies the transport type of the address. ipv4Tdomain specifies the transport type of address is an IPv4 address and ipv6Tdomain specifies the transport type of address is IPv6. The default is ipv4Tdomain.
TAddress	Specifies the transport address in xx.xx.xx.xx:port format, for example: 10:10:10:10:162, where 162 is the trap listening port on the system 10.10.10.10.
Timeout	Specifies the maximum round trip time required to communicate with the transport address. The value is in 1/100 seconds from 0–2147483647. The default is 1500. After the system sends a message to this address, if a response (if one is expected) is not received within this time period, you can assume that the response is not delivered.
RetryCount	Specifies the maximum number of retries if a response is not received for a generated message. The count can be in the range of 0–255. The default is 3.
TagList	Contains a list of tag values used to select target addresses for a particular operation. A tag refers to a class of targets to which the messages can be sent.
Params	Contains SNMP parameters used to generate messages to send to this transport address. For example, to receive SNMPv2C traps, use TparamV2.

Table continues...

Name	Description
TMask	Specifies the mask. The value can be empty or in six-byte hex string format. Tmask is an optional parameter that permits an entry in the TargetAddrTable to specify multiple addresses.
MMS	Specifies the maximum message size. The size can be zero, or 484–2147483647. The default is 484. Although the maximum MMS is 2147483647, the device supports the maximum SNMP packet size of 8192.

Configuring target table parameters

Configure the target table to configure the security parameters for SNMP. Configure the target table to configure parameters such as SNMP version and security levels.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > SnmpV3**.
2. Click **Target Table**.
3. Click the **Target Params Table** tab.
4. Click **Insert**.
5. In the **Name** box, type a target table name.
6. From the **MPModel** options, select an SNMP version.
7. From the **Security Model** options, select the security model.
8. In the **SecurityName** box, type `readview` or `writeview`.
9. From the **SecurityLevel** options, select the security level for the table.
10. Click **Insert**.

Target Params Table field descriptions

Use the data in the following table to use the **Target Params Table** tab.

Name	Description
Name	Identifies the target table.
MPModel	Specifies the message processing model to use to generate messages: SNMPv1, SNMPv2c, or SNMPv3/USM.
SecurityModel	Specifies the security model to use to generate messages: SNMPv1, SNMPv2c, or USM. You can receive an inconsistentValue error if you try to configure this variable to a value for a security model that the implementation does not support.

Table continues...

Name	Description
SecurityName	Identifies the principal on whose behalf SNMP messages are generated.
SecurityLevel	Specifies the security level used to generate SNMP messages: noAuthNoPriv, authNoPriv, or authPriv.

Configuring an SNMP notify table

Configure the notify table to select management targets to receive notifications, as well as the type of notification to send to each management target.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > SnmpV3**.
2. Click **Notify Table**.
3. In the **Notify Table** tab, click **Insert**.
4. In the **Name** box, type a notify table name.
5. In the **Tag** box, type the transport tag for the table.
6. From the **Type** options, select a type.
7. Click **Insert**.

Notify Table field descriptions

Use the data in the following table to use the **Notify Table** tab.

Name	Description
Name	Specifies a unique identifier.
Tag	Specifies the tag.
Type	<p>Determines the type of notification generated. This value is only used to generate notifications, and is ignored for other purposes. If an SNMP entity only supports generation of Unconfirmed-Class protocol data unit (PDU), this parameter can be read-only. The possible values are</p> <ul style="list-style-type: none"> • trap—Messages generated contain Unconfirmed-Class Protocol Data Units (PDU). • inform—Messages generated contain Confirmed-Class PDUs. <p>The default value is trap.</p>

Configuring SNMP notify filter profiles

Configure the SNMP table of filter profiles to determine whether particular management targets receive particular notifications.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > SnmpV3**.
2. Click **Notify Table**.
3. Click the **Notify Filter Table** tab.
4. Click **Insert**.
5. In the **NotifyFilterProfileName** box, type a name for the notify filter profile.
6. In the **Subtree** box, type subtree location information in x.x.x.x.x.x.x.x. format.
7. In the **Mask** box, type the mask location in hex string format.
8. From the **Type** options, select **included** or **excluded**.
9. Click **Insert**.

Notify Filter Table field descriptions

Use the data in the following table to use the **Notify Filter Table** tab.

Name	Description
NotifyFilterProfileName	Specifies the name of the filter profile used to generate notifications.
Subtree	Specifies the MIB subtree that, if you combine it with the mask, defines a family of subtrees, which are included in or excluded from the filter profile. For more information, see RFC2573.
Mask	Specifies the bit mask (in hexadecimal format) that, in combination with Subtree, defines a family of subtrees, which are included in or excluded from the filter profile.
Type	Indicates whether the family of filter subtrees are included in or excluded from a filter. The default is included.

Configuring SNMP notify filter profile table parameters

Configure the profile table to associate a notification filter profile with a particular set of target parameters.

Before you begin

- The notify filter profile exists.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > SnmpV3**.

2. Click **Notify Table**.
3. Click the **Notify Filter Profile Table** tab.
4. Click **Insert**.
5. In the **TargetParamsName** box, type a name for the target parameters.
6. In the **NotifyFilterProfileName** box, type a name for the notify filter profile.
7. Click **Insert**.

Notify Filter Profile Table field descriptions

Use the data in the following table to use the **Notify Filter Profile Table** tab.

Name	Description
TargetParamsName	Specifies the unique identifier associated with this entry.
NotifyFilterProfileName	Specifies the name of the filter profile to use to generate notifications.

Enabling SNMP trap logging

Enable trap logging to save a copy of all SNMP traps.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **General**.
3. Click the **Error** tab.
4. Select **AuthenticationTraps**.
5. Click **Apply**.

Error field descriptions

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

Name	Description
AuthenticationTraps	Enables or disables the sending of traps after an error occurs. The default is disabled.
LastErrorCode	Specifies the last reported error code.
LastErrorSeverity	Specifies the last reported error severity: 0= Informative Information 1= Warning Condition 2= Error Condition 3= Manufacturing Information 4= Fatal Condition

Viewing the trap sender table

Use the Trap Sender Table tab to view source and receiving addresses.

Procedure

1. On the Device Physical View, select a chassis.
2. In the navigation pane, expand the following folders: **Configuration > Edit**.
3. Click **Chassis**.
4. Click the **Trap Sender Table** tab.

Trap Sender Table field descriptions

Use the data in the following table to use the **Trap Sender Table** tab.

Name	Description
RecvAddress	IP address for the trap receiver. This is a read-only parameter that contains the IP address configured in the TAddress field in the TargetTable.
SrcAddress	Source IP address to use when sending traps. This IP address will be inserted into the source IP address field in the UDP trap packet.

Chapter 6: Hardware troubleshooting

The following sections provide troubleshooting information for common hardware problems.

Troubleshooting module failure

If a module failure occurs, check for possible midplane connection problems. Make sure that you correctly seated the module in the midplane connector and that you securely tightened the retaining screws.

If a module fails during module initialization and the replacement module is the same module type, in rare cases, the new module does not initialize.

Procedure

1. Remove the faulty module.
2. Insert a new module.
3. Restart the chassis.

If the new module fails to initialize, perform the following procedure steps.

4. Remove the faulty module.
5. Insert a different module type from the module type removed in Step 1, and then wait for the replacement module to initialize.

 **Caution:**

Before you insert a different module type, save the configuration. Do not save the configuration during the testing phase or you will lose the configuration for that module.

6. Remove the module inserted in the preceding step.
7. Insert the new module in the slot where the faulty module resided. This new module model must be identical to the model removed in Step 1.

If the module still fails to operate, contact the Avaya Technical Solutions Center for assistance.

Troubleshooting CP start failure

A troubleshooting menu appears in instances where a CP module fails to start. Depending on the type of failure detected, you see one of two recovery menus. The following example illustrates the first menu.

```

HW faults detected:
** Internal Compact Flash not mounted

*****
*
* WARNING:
* The Lifecycle recovery options are used to recover the /intflash
* in the events of corruption as well as resetting the login/password
* back to default (rwa/rwa). Data that includes configuration files,
* log files, core files, etc. stored originally on the /intflash
* could be lost in the recovery attempt.
*
*****

Lifecycle recovery menu

1 - Save all config files in /intflash but not subdirectories;
  Save primary and secondary software releases;
  Reformat /intflash; Reboot
2 - Reformat /intflash; Reboot

q - Quit

Please make your selection:

```

The following example illustrates the second menu.

```

HW faults detected:
** Internal Compact Flash not detected

*****
*
* WARNING:
* A hardware failure has been detected that prevents proper operation
* of this module. You can attempt a system reboot to try to recover
* from the failure. If a reboot does not correct the problem, then
* leave the module hung at this prompt and contact customer support.
*
*****

Lifecycle HW recovery menu

1 - Reboot system to attempt recovery;

Please make your selection:

```

Removing external storage devices from the CP module

Perform this procedure to safely remove the USB and the external Compact Flash devices from the CP module. You must perform this procedure to prevent data loss or hardware damage.

! Important:

Do not unplug the storage device without first performing this procedure.

You must use the appropriate stop command to unmount the device before you physically remove it from the CP module.

Before you begin

Several system tools use the external Compact Flash as the default storage location. Check the following features before you remove the card:

- Packet Capture (PCAP)
- logging
- debug or trace

The Virtual Services Platform 9000 stop command does not succeed if the specified device is in use. Common uses that impede the proper execution of the stop command are:

- USB or external Compact Flash file access is in progress (move, copy, read, or write) to or from the USB, or the external Compact Flash.

Discontinue operations or wait for access completion before you use the stop command.

- The ACLI session current working directory is configured for the device you need to remove. Change the current working directory to internal Compact Flash, which is the default.
- Logging is enabled to the external Compact Flash, which is the default.

Use the `show logging config` command to verify the current storage location. If the location is the external Compact Flash card that you need to remove, use the `no logging logToExtFlash` command to log to the internal Compact Flash.

- PCAP is enabled.

Disable PCAP, which requires the external Compact Flash. Use the `show pcap` command to verify if PCAP is enabled. To disable PCAP, use the `no pcap enable` command.

- Debugging features are enabled.

The debug-config file and trace-logging flags must be disabled, which is the default. Use the `show boot config flags` command to verify the status. Use the `no boot config flags debug-config file` or the `no boot config flags trace-logging` command to disable these flags.

About this task

*** Note:**

Use the Avaya Compact Flash device (EC1411010-E6) with the Virtual Services Platform 9000 because the Avaya Compact Flash is validated for proper operation on the Virtual Services Platform 9000. Do not use other Compact Flash devices because they are not verified for Virtual Services Platform 9000 compatibility, and can result in loss of access to the Compact Flash device.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```

2. Remove a USB device:

a. Unmount the USB device:

```
usb-stop
```

b. Wait for the response that indicates it is safe to remove the device.

c. Physically remove the device.

3. Remove an external Compact Flash device:

a. Unmount the external flash device:

```
extflash-stop
```

b. Wait for the response that indicates it is safe to remove the device.

c. Physically remove the device.

Example

Unmount and remove the USB:

```
VSP-9012:1>enable
VSP-9012:1#usb-stop
It is now safe to remove the USB device.
VSP-9012:1#extflash-stop
It is now safe to remove the external Compact Flash device.
```

Next steps

No restrictions or requirements exist before you can reinsert a USB or external Compact Flash device. You can insert these devices at any time and Virtual Services Platform 9000 automatically recognizes them. The devices are accessible within seconds after insertion.

After you insert the external Compact Flash, enable logging to the external Compact Flash with the `logging logToExtFlash` command.

Additionally, you can enable the following features as required:

- PCAP
- debug-config file or trace-logging flags

Troubleshooting USB viewing problems

After you insert a USB device in the USB slot, the Linux system automatically detects and mounts the device. If you cannot view files on the device, perform this procedure.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Check the file system:

```
ls /usb/
```

3. Remove a USB device:

a. Unmount the USB device:

```
usb-stop
```

b. Wait for the response that indicates it is safe to remove the device.

c. Physically remove the device.

4. Remove and then reinsert the device.

5. Check the device for errors:

```
dos-chkdsk /usb
```

6. If errors are detected, then you can reformat the device:

```
dos-format /usb
```

Caution:

If you format the device, you erase all data on the device.

Example

Check the file system:

```
VSP-9012:1>enable
VSP-9012:1#ls /usb/
Listing Directory /usb/:
drwxr-xr-x 4 0 0 4096 Jan 1 1970 ./
drwxrwxr-x
22 0 0 0 Sep 9 20:22 ../
drwxr-xr-x 2 0 0 4096 Mar 17 16:03 Photos-of-Flash-
drwxr-xr-x 2 0 0 4096 Jun 13 20:56 intflash/
```

Check the device for errors:

```
VSP-9012:1#usb-stop
It is now safe to remove the USB device.
VSP-9102:1#dos-chkdsk /usb
/usr/sbin/fsck.vfat /dev/usb1 -v >& /dev/console dosfsck 2.11a
(05 Mar 2010)
dosfsck 2.11a, 05 Mar 2010, FAT32, LFN
Checking we can access the last sector of the filesystem
Boot sector contents:
System ID "mkdosfs"
Media byte 0xf8 (hard disk)
512 bytes per logical sector
4096 bytes per cluster
32 reserved sectors
First FAT starts at byte 16384 (sector 32)
2 FATs, 32 bit entries
3897344 bytes per FAT (= 7612 sectors)
Root directory start at cluster 2 (arbitrary size)
Data area starts at byte 7811072 (sector 15256)
974240 data clusters (3990487040 bytes)
62 sectors/track, 124 heads
0 hidden sectors
7809178 sectors total
Checking for unused clusters.
```

```
Checking free cluster summary.  
/dev/usb1: 17 files, 174804/974240 clusters
```

If errors are detected, reformat the disk:

```
VSP-9012:1#dos-format /usb
```

Troubleshooting USB writing problems

Use the following information to troubleshoot USB writing problems.

About this task

USB storage devices typically provide a switch to write-protect the device data; the location and movement of this switch depends on the vendor and device.

A write problem can also occur when a write operation indicates that the device has no more room, but a directory listing of the device shows considerable free space available.

In this case, the device is improperly formatted with a FAT-16 file system, which limits the number of files in the root directory to 256.

Procedure

1. Verify that the write-protect switch is in the write position.
2. To check if you properly formatted the device, try to create file number 257. You know you improperly formatted the device if the device gives the no more space error.
3. You can delete files from the device.
or
4. Reformat the device. See Troubleshooting USB Viewing Problems for more information: [Troubleshooting USB viewing problems](#) on page 97.

Troubleshooting external Compact Flash viewing problems

After you insert an external Compact Flash in the Compact Flash slot, the Linux system automatically detects and mounts the device. If you cannot view files on the device, perform this procedure.

Procedure

1. Enter Global Configuration mode:
enable
configure terminal

2. Check the file system:

```
ls /extflash/
```

3. Remove an external Compact Flash:

a. Unmount the external Compact Flash:

```
extflash-stop
```

b. Wait for the response that indicates it is safe to remove the Compact Flash.

c. Physically remove the Compact Flash.

4. Remove and then reinsert the device.

5. Check the device for errors:

```
dos-chkdisk /extflash
```

6. If errors are detected, then you can reformat the Compact Flash:

```
dos-format /extflash
```

 **Caution:**

If you format the device, you erase all data on the device.

Example

Check the file system. Unmount the external Compact Flash. Check the device for errors, and reformat if errors are detected.

```
VSP-9012:1>enable
VSP-9012:1#ls /extflash/
VSP-9012:1#extflash-stop
It is now safe to remove the external Compact Flash device.
VSP-9012:1#dos-chkdisk /extflash
VSP-9012:1#dos-format /extflash
```

Using trace to diagnose hardware problems

About this task

Use trace to observe the status of a hardware module at a certain time.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Begin the trace operation:

```
line-card {<3-12>|SF1|SF2|SF3|SF4|SF5|SF6} trace level [{67-179} {0-4}]
```

3. Search the trace for a specific string value:

```
line-card {<3-12>|SF1|SF2|SF3|SF4|SF5|SF6} trace grep [WORD<0-1024>]
```

Example

Begin the trace operation. Search the trace for a specific string value.

```
VSP-9012:1>enable
VSP-9012:1#line-card SF1 trace level 67 1
VSP-9012:1#line-card SF1 trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

Use the data in the following table to use the `line-card` command.

Table 9: Variable definitions

Variable	Value
{<3-12> SF1 SF2 SF3 SF4 SF5 SF6}	Specifies the slot number for the interface module or Switch Fabric module.
{67-179} {0-4}	Starts the trace by specifying the module ID and level. <67-179> specifies the module ID. <0-4> specifies the trace level from 0-4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
<i>WORD</i> <0-1024>	Performs a string search in the trace.

Chapter 7: Software troubleshooting

Use the following sections to find general troubleshooting information for Avaya Virtual Services Platform 9000.

Software troubleshooting

This section contains general troubleshooting information for Avaya Virtual Services Platform 9000 software.

Failure to read the configuration file

The device can fail to read and load a saved configuration file after it starts. This situation occurs if you enable the `factorydefaults` boot configuration flag. Configure the flag to false: `no boot config flags factorydefaults`.

Example

Configure the flag to false:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#no boot config flags factorydefaults
```

No web management interface access to a device

If the device and the PC that runs the web browser are in the same network, you can find that even though other applications, for example, Telnet, can access a particular switch, the web management interface cannot. This situation can occur if the web browser has a proxy server that resolves the `www` path and returns the reachable IP address to the browser. If no route exists from the proxy server to the device, the HTTP query does not reach the device, and does not receive a response.

To prevent this problem, ensure that if the web browser uses a proxy server, you specify a route from the proxy server to the device.

Cannot enable encryption

To enable encryption on the system, you must first download and install the necessary encryption module. If you do not download and install the module, you can enable the encryption method in the software but it does not work. You must download and install the encryption modules separately from the software releases. For more information about how to download and install an encryption module, see *Applying Upgrades and Patches to Avaya Virtual Services Platform 9000*, NN46250-400.

Debug files

Virtual Services Platform 9000 stores debug files in the intflash directory.

The debug file is in a zipped format and contains information to help debug the device, including:

- a memory snapshot
- logs
- traces

Avaya recommends you delete these files to ensure enough space exists in the internal flash on the CP module. New files do not overwrite old files. You must remove the files; otherwise, the internal flash may not have enough free space for necessary activities, for example, to store a core dump file if the switch fails, or you may not have the space to transfer a new release to the internal flash to upgrade your switch.

Virtual Services Platform 9000 stores a maximum of 32 files for each debug file for each slot, depending on the file size of each debug file. The internal flash provides 2 GB of storage. A message appears on the console to inform you when less than 700 MB is available.

The `debug-file remove` command can delete the following types of debug files:

- core
- archive
- PMEM
- dmalloc
- flrec
- wd_stats

If you want to delete a specific file, you must use the `remove` command. For more information, see *ACLI Commands Reference for Avaya Virtual Services Platform 9000*, NN46250-104.

High Availability-CPU

VSP 9000 supports High Availability-CPU mode for the `show debug-file` and `debug-file remove` commands.

SNMP

Virtual Services Platform 9000 does not support SNMP for the `show debug-file` or the `debug-file remove` commands.

Software download

This section describes where to download software or documentation.

Downloading the software

Download new software to upgrade the Avaya Virtual Services Platform 9000. Software downloads can include encryption modules and software images.

Download patches and readme files from the Avaya support site at www.avaya.com/support.

Before you begin

- You must have access to the new software from the Avaya support site: www.avaya.com/support. You need a valid user or site ID and password.

About this task

Download the Advanced Encryption Standard (AES) and Data Encryption Standard (DES) software before you enable the encryption algorithms and use SNMPv3. The AES and DES encryption modules exist in a single file and you can enable them when the file is stored on flash.

Download the SSH encryption software before you enable the 3DES encryption module and use SSH.

Due to export restrictions, the encryption capability is separate from the main software image. SNMPv3 and the SSH server do not function properly without the use of this image.

For more information about file names for the current release, see *Release Notes for Avaya Virtual Services Platform 9000*, NN46250-401.

Important:

You must load the security encryption modules on the device before you can use the protocol.

Procedure

1. From an Internet browser, browse to www.avaya.com/support.
2. Click **DOWNLOADS**.
3. In the product search field, type **Virtual Services Platform 9000**.
4. In the **Choose Release** field, click a release number.
5. Click the download title to view the selected information.
6. Click the file you want to download.

7. Login to download the required software file.
8. Use an FTP client in binary mode to transfer the file to either the Virtual Services Platform 9000 or an external USB device.

Downloading Avaya Virtual Services Platform 9000 documentation

About this task

Download documentation from the Avaya website to obtain conceptual, procedural, and referential information for the Virtual Services Platform 9000.

Procedure

1. From an Internet browser, browse to www.avaya.com/support.
2. Click **Products**.
3. Enter **Virtual Services Platform 9000**, and then select the link that appears.
4. Click **Documentation**.
5. Click a job function or **View All Documents**.

You can filter the document list by release number.

Software troubleshooting tool configuration using the ACLI

Use the tools described in this section to perform troubleshooting procedures using ACLI.

Using ACLI for troubleshooting

You can use ACLI to provide diagnostic information.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Disable scrolling of the output display:
`terminal more disable`
3. View configuration file information:
`more WORD<1-99>`

- Capture the following output, which displays the current switch configuration, after you observe a problem with the device:

```
show running-config [verbose] [module <boot|cfm|cli|cluster|diag|
filter|ip|ipv6|isis|l2|l3|lacp|mlt|naap|nsna|ntp|port|qos|radius|rmon|
slamon|slpp|spbm|stg|sys|tacacs|vlan|vsptalk|web>]
```

- Capture the following output after you observe a problem with the device. The command displays technical information about the status of the system and complete information about the hardware components, software components, and operation of the system:

```
show tech
```

- Capture the following output after you observe a problem with the device. The command displays individual statistics to manage network performance:

```
show interfaces gigabitEthernet statistics <bridging {slot/port[-
slot/port][,...]}|dhcp-relay {slot/port[-slot/port][,...]} [vrf
WORD<0-16>][vrfids WORD<0-512>]|lacp {slot/port[-slot/port][,...]}|
policer {slot/port[-slot/port][,...]}|rmon {slot/port[-slot/port]
[,...]}[history]|verbose {slot/port[-slot/port][,...]}>
```

- Capture the following output after you observe a problem with the device. The command displays general error information for the port:

```
show interfaces gigabitEthernet error <collision|ospf|verbose>
{slot/port[-slot/port][,...]}
```

Example

Capture the following output after you observe a problem with the device. The command displays the current switch configuration.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#show running-config module cli
Preparing to Display Configuration...
#
# Sun Jun 23 19:31:50 2013 EDT
# box type           : VSP-9012
# software version   : 3.4.0.0
# cli mode           : ACLI
#
#ASIC Info :
#Slot #1:
#   Module: 9080CP
#   OXATE CPLD: 10032310
#   OXIDE FPGA: 12041711
#   CATSKILL FPGA: 13041115
#   QE version: QE2000_A0
#Slot #4:
#   Module: 9048GB
#   K2 FPGA: 12030509
#   IODATEDC CPLD: 09041015
#   IODATEBB CPLD: 09041016
#   PIM48SFP CPLD: 09050110
#   SULFIDE FPGA: 10041310
```

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

Capture the following output after you observe a problem with the device. The command displays technical information about the status of the system and the complete information about the hardware components, software components, and operation of the system.

```
VSP-9012:1(config)#show tech
```

```
Sys Info:
```

```
-----
```

```
General Info :
```

```

SysDescr      : VSP-9012 (4.0.0.0) (GA)
SysName       : VSPROF
SysUpTime     : 0 day(s), 00:32:15
SysContact    : http://support.avaya.com/
SysLocation   : 211 Mt. Airy Road,Basking Ridge,NJ 07920

```

```
Chassis Info:
```

```

Chassis       : 9012
Serial#       : LBNNTMRJ0000DF
H/W Revision  : 323500-A 01
H/W Config    : 0100
NumSlots      : 12
NumPorts     : 122
BaseMacAddr   : 80:17:7d:75:00:00
MacAddrCapacity : 4096

```

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

Capture the following output after you observe a problem with the device. The command displays individual statistics for specific ports to manage network performance.

```
VSP-9012:1#show interfaces gigabitethernet statistics
```

```
=====
```

Port Stats Interface					
PORT NUM	IN OCTETS	OUT OCTETS	IN PACKET	OUT PACKET	
4/1	1215232	1852156	18988	25083	
4/2	11866260	3650340	128847	51849	
4/3	0	0	0	0	
4/4	0	0	0	0	
4/5	0	0	0	0	
4/6	2606433776	2605569408	40718802	40712022	
4/7	2383797478	2368788480	37189478	37012320	
4/8	2639779622	2624836140	41201664	40945760	
4/9	0	0	0	0	
4/10	0	0	0	0	
4/11	0	0	0	0	
4/12	0	6776546	0	62572	
4/13	1215232	997632	18988	15588	
4/14	7459408	1396224	69625	18702	

```
-----
```

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

Capture the following output after you observe a problem with the device. The command displays general error information for the port.

```
VSP-9012:1#show interfaces gigabitEthernet error
```

```
=====
                        Port Ethernet Error
=====
PORT  ERROR  ERROR  FRAMES  TOO    LINK   CARRIER  CARRIER  SQETEST  IN
NUM   ALIGN  FCS    LONG    SHORT  FAILURE SENSE     ERRORS    ERRORS   DISCARD
-----
4/1   0       0      0       0      0      0         0         0         0
4/2   0       0      0       0      0      0         0         0         0
4/3   0       0      0       0      0      0         0         0         0
4/4   0       0      0       0      0      0         0         0         0
4/5   0       0      0       0      0      0         0         0         0
4/6   0       0      0       0      0      0         0         0         0
4/7   0       0      0       0      0      0         0         0         0
4/8   0       0      0       0      0      0         0         0         0
4/9   0       0      0       0      0      0         0         0         0
4/10  0       0      0       0      0      0         0         0         0
4/11  0       0      0       0      0      0         0         0         0
4/12  0       0      0       0      0      0         0         0         0
4/13  0       0      0       0      0      0         0         0         0
4/14  0       0      0       0      0      0         0         0         0
4/15  0       0      0       0      0      0         0         0         0
4/16  0       0      0       0      0      0         0         0         0

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

Variable definitions

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

Table 10: Variable definitions

Variable	Value
<i>WORD</i> <1–99>	Specifies the file name to view. Provide the filename in one of the following formats: a.b.c.d:<file>, x:x:x:x:x:x:<file>, /intflash/<file>, /extflash/<file>, or /usb/<file>.

Use the data in the following table to use the **show running-config** command.

Table 11: Variable definitions

Variable	Value
module <boot cfm cli cluster diag filter ip ipv6 isis lcp mlt naap nsna ntp port qos radius rmon slamon slpp spbm stg sys tacacs vlan vsptalk web>	Specifies the command group for which you request configuration settings.
verbose	Specifies a complete list of all configuration information about the switch.

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

Variable	Value
bridging {slot/port[-slot/port][,...]}	Displays ports bridging statistics.
dhcp-relay {slot/port[-slot/port][,...]} [vrf WORD<0–16>][vrfs WORD<0–512>]	Displays port Dynamic Host Configuration Protocol (DHCP) statistics.
lacp {slot/port[-slot/port][,...]}	Displays Link Aggregation Control Protocol (LACP) statistics.
policer {slot/port[-slot/port][,...]}	Displays policer statistics.
rmon {slot/port[-slot/port][,...]}[history]	Displays Remote Network Monitoring (RMON) statistics.
verbose {slot/port[-slot/port][,...]}	Displays a complete list of all statistics.

Use the data in the following table to use the **show interfaces gigabitEthernet error** command.

Table 12: Variable definitions

Variable	Value
collision	Displays port collision error information.
ospf	Displays port Open Shortest Path First (OSPF) error information.
verbose	Displays all port error information.
{slot/port[-slot/port][,...]}	Specifies the port.

Using software record dumps

About this task

Capture a dump of the software records from ingress traffic to help troubleshoot performance problems. Generally, a verbosity level of 1 suffices.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Dump software record information:


```
dump ar <1-12> WORD<1-1536> <0-3>
```

Example

Dump software record information:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#dump ar 1 vlan 1
```

Variable definitions

Use the data in the following table to use the `dump ar` command.

Table 13: Variable definitions

Variable	Value
<1-12>	Specifies the slot number.
WORD<1-1536>	Specifies a record type in the AR table. Options include vlan, ip_subnet, mac_vlan, mac, arp, ip, ipmc, ip_filter, protocol, all.
<0-3>	Specifies the verbosity from 0–3. Higher numbers specify more verbosity.

Using trace to diagnose problems

Use trace to observe the status of a software module at a certain time.

Before you begin

 **Caution:**

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

About this task

If you notice a CPU utilization issue (generally a sustained spike above 90%) perform a trace of the control plane activity.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Clear the trace:
`clear trace`
3. Identify the module ID for which you want to use the trace tool:
`show trace modid-list`
4. Begin the trace operation:
`trace level [<0-192>] [<0-4>]`
5. Wait approximately 30 seconds.

The default trace settings for CPU utilization are:

- High CPU Utilization: 90%
- High Track Duration: 5 seconds

- Low CPU Utilization: 75%
- Low Track Duration: 5 seconds

6. Stop tracing:

```
trace shutdown
```

7. View the trace results:

```
show trace file [tail]
```

8. Save the trace file to the Compact Flash card for retrieval.

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

9. Search trace results for a specific string value, for example, the word error:

```
trace grep [WORD<0-128>]
```

If you use this command and do not specify a string value, you clear the results of a previous search.

10. Stop tracing:

```
trace shutdown
```

Example

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

Clear the trace:

```
VSP-9012:1#clear trace
```

Identify the module ID for which you want to use the trace tool:

```
VSP-9012:1#show trace modid-list
```

```
0 - COMMON
1 - SNMP
2 - RMON
3 - PORT_MGR
4 - CHAS_MGR
5 - BRIDGE
6 - OSPF
7 - HWIF
8 - SIM
9 - CPP
10 - NETDRV
11 - VLAN_MGR
12 - CLI
13 - MAIN
14 - P2IP
15 - RCIP
16 - WEBSRV
17 - ACIF
18 - GBIF
20 - TDP
21 - MAN_DIAG
```

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

Begin the trace operation:

```
VSP-9012:1#trace level 2 3
```

Stop tracing:

```
VSP-9012:1#trace shutdown
```

Save the trace file to the Compact Flash card for retrieval:

```
VSP-9012:1#save trace
```

Search trace results for a specific string value, for example, the word error:

```
VSP-9012:1#trace grep error
```

Search trace results for a specific string value, for example, MAC address 00-1A-4B-8A-FB-6B:

```
VSP-9012:1#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

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

Table 14: Variable definitions

Variable	Value
grep [WORD<0-128>]	Search trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level [<0-192>] [<0-4>]	Starts the trace by specifying the module ID and level. <ul style="list-style-type: none"> • <0-192> specifies the module ID. • <0-4> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
shutdown	Stops the trace operation.
screen {disable enable}	Enables the display of trace output to the screen.

Use the data in the following table to use the **save trace** command.

Table 15: Variable definitions

Variable	Value
file WORD<1-99>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d:<file> • x:x:x:x:x:x:x <file> • /intflash/<file> • /extflash/<file>

Variable	Value
	<ul style="list-style-type: none"> • /usb/<file> • /mnt/intflash/ <file> • /mnt/extflash/ <file> <p>/mnt/intflash is the internal flash of the second CP module (the one to which you are not connected) .</p> <p>/mnt/extflash is the external flash of the second CP module (the one to which you are not connected) .</p>

Using trace to diagnose Ipv6 problems

Before you begin

Caution:

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

About this task

Use trace to observe the status of IPv6 at a certain time.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Activate or deactivate the trace for the IPv6 base:

```
trace ipv6 base <disable|enable> <all|debug|error|icmp|info|
ipclient|nbr|pkt|warn>
```

3. Activate or deactivate the trace for IPv6 forwarding:

```
trace ipv6 forwarding <disable|enable> <all|debug|error|info|pkt|
warn>
```

4. Activate or deactivate the trace for IPv6 neighbor discovery:

```
trace ipv6 nd <disable|enable> <all|debug|error|info|nbr|pkt|
redirect|warn>
```

5. Activate or deactivate the trace for IPv6 OSPF:

```
trace ipv6 ospf <disable|enable> <adj|all|config|error|import|info|
lsa|pkt|spf|warn>
```

6. Activate or deactivate the trace for the IPv6 routing table manager:

```
trace ipv6 rtm <disable|enable> <all|change-list|debug|error|fib|
info|redist|update|warn>
```

7. Activate or deactivate the trace for IPv6 transport:

```
trace ipv6 transport <disable|enable> <all|common|tcp|udp>
```

Example

```
Switch:1>enable
```

Activate the trace for all the IPv6 base categories:

```
Switch:1#trace ipv6 base enable all
```

Activate the trace for all the IPv6 forwarding categories:

```
Switch:1#trace ipv6 forwarding enable all
```

Activate the trace for all the IPv6 neighbor discovery categories:

```
Switch:1#trace ipv6 nd enable all
```

Activate the trace for the all IPv6 routing table manager categories:

```
Switch:1#trace ipv6 rtm enable all
```

Activate the trace for all the IPv6 transport categories:

```
Switch:1#trace ipv6 transport enable all
```

Variable definitions

Use the data in the following table to use the **trace ipv6** command.

Table 16: Variable definitions

Variable	Value
base <disable enable> <all debug error icmp info ipclient nbr pkt warn>	Enables or disables a specific trace category for IPv6 base.
forwarding <disable enable> <all debug error info pkt warn>	Enables or disables a specific trace category for IPv6 forwarding.
nd <disable enable> <all debug error info nbr pkt redirect warn>	Enables or disables a specific trace category for IPv6 neighbor discovery.
ospf <disable enable> <adj all config error import info lsa pkt spf warn>	Enables or disables a specific trace category for IPv6 OSPF.
rtm <disable enable> <all change-list debug error fib info redist update warn>	Enables or disables a specific trace category for IPv6 routing table manager.
transport <disable enable> <all common tcp udp>	Enables or disables a specific trace category for IPv6 transport.

Viewing and clearing the alarm database

View the alarm database regularly to monitor alarm conditions, even if you do not observe a performance problem. Review the alarm messages to determine if the system performs as expected.

Not all alarm conditions indicate a problem so you must be familiar with expected behavior.

Procedure

1. Log on to the switch to enter User EXEC mode.

2. Display the contents of the alarm database:

```
show alarm database
```

The alarm ID, severity, creation time, and reason fields are the most important fields to use when isolating a fault.

3. **(Optional)** View the alarm database statistics:

```
show alarm statistics
```

4. Log on to Privileged EXEC mode:

```
enable
```

5. **(Optional)** Clear a specific database event by the alarm ID:

```
clear alarm database alarm-id WORD<0-100>
```

6. **(Optional)** Clear all alarms in the database:

```
clear alarm database
```

7. **(Optional)** Clear alarm statistics:

```
clear alarm statistics
```

Example

View the alarm database and the alarm statistics.

```
VSP-9012:1>show alarm database
```

ALARM ID	EVENT CODE	ALARM TYPE	ALARM STATUS	SEVERITY	FREQ	CREATION TIME	UPDATED TIME	CLEARED TIME	REASON
00000005.1	0x00000661	DYNAMIC	SET	WARNING	1	[06/29/12 11:03:56.854]	[06/29/12 11:03:56.854]	[--/-- --:--:--]	Slot 1: Logging to internal flash is not recommended - please insert external flash and ensure logging to external flash is con
0040000a.2	0x0001072b	DYNAMIC	SET	WARNING	1	[06/29/12 11:04:04.632]	[06/29/12 11:04:04.632]	[--/-- --:--:--]	Unsupported Power Supply Detected in slot PS 2.
0040000a.5	0x0001072b	DYNAMIC	SET	WARNING	1	[06/29/12 11:04:04.633]	[06/29/12 11:04:04.633]	[--/-- --:--:--]	Unsupported Power Supply Detected in slot PS 5.
00400006.3	0x000106cc	DYNAMIC	SET	WARNING	1	[06/29/12 11:04:04.634]	[06/29/12 11:04:04.634]	[--/-- --:--:--]	No fan module is present in slot SF-FAN 1
00300001.258	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.003]	[06/29/12 11:04:52.003]	[--/-- --:--:--]	Link Down(4/3)
00300001.260	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.004]	[06/29/12 11:04:52.004]	[--/-- --:--:--]	Link Down(4/5)
00300001.261	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.006]	[06/29/12 11:04:52.006]	[--/-- --:--:--]	Link Down(4/6)
00300001.262	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.007]	[06/29/12 11:04:52.007]	[--/-- --:--:--]	Link Down(4/7)
00300001.263	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.008]	[06/29/12 11:04:52.008]	[--/-- --:--:--]	Link Down(4/8)
00300001.267	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.010]	[06/29/12 11:04:52.010]	[--/-- --:--:--]	Link Down(4/12)
00300001.278	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.014]	[06/29/12 11:04:52.014]	[--/-- --:--:--]	Link Down(4/23)
00300001.279	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.015]	[06/29/12 11:04:52.015]	[--/-- --:--:--]	Link Down(4/24)
00300001.282	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.019]	[06/29/12 11:04:52.019]	[--/-- --:--:--]	Link Down(4/27)
00300001.283	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.020]	[06/29/12 11:04:52.020]	[--/-- --:--:--]	Link Down(4/28)
00300001.300	0x0000c5e7	DYNAMIC	SET	INFO	1	[06/29/12 11:04:52.025]	[06/29/12 11:04:52.025]	[--/-- --:--:--]	Link Down(4/45)
09200003.1	0x00248546	DYNAMIC	SET	WARNING	1	[06/29/12 11:04:54.838]	[06/29/12 11:04:54.838]	[--/-- --:--:--]	SYSTEM HAS 1 BME - No Switch Fabric Redundancy Control !!!
00400005	0x000045e5	DYNAMIC	SET	INFO	1	[06/29/12 11:05:38.852]	[06/29/12 11:05:38.852]	[--/-- --:--:--]	Sending Cold-Start Trap
00000006.1	0x000005a7	DYNAMIC	SET	WARNING	1	[06/29/12 11:05:38.853]	[06/29/12 11:05:38.853]	[--/-- --:--:--]	No configured hosts are reachable for log file transfer

```
VSP-9012:1>show alarm statistics
```

```
=====
                          ALARM STATISTICS
=====
PERSISTENT PERSISTENT PERSISTENT PERSISTENT DYNAMIC DYNAMIC DYNAMIC DYNAMIC
ALARM      ACTIVE      CLEARED      WRPRD      ALARM      ACTIVE      CLEARED      WRPRD
0           0           0             0           34         18         16           0
```

Viewing and deleting debug files

Use this procedure to view and delete debug files.

Delete debug files to free space in the intflash, which has 2 GB of space. Avaya recommends you delete these files to ensure enough space exists in intflash. New debug files do not overwrite old debug files. You must remove the file; otherwise, enough free space may not exist in the intflash to store the core dump if the switch fails or enough space may not exist for you to transfer a new release to the intflash of the switch to upgrade your switch.

The `debug-file remove` command can delete the following types of files:

- core
- archive
- PMEM
- dmalloc
- flrec
- wd_stats

If you want to delete a specific file, you must use the `remove` command. For more information, see *ACLI Commands Reference for Avaya Virtual Services Platform 9000*, NN46250-104.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View debug files:

```
show debug-file [all][{slot[-slot][,....]}
```

3. Delete debug files:

```
debug-file remove [all][{slot[-slot][,....]}
```

4. Enter Privileged EXEC mode:

```
enable
```

5. View core files:

```
show core-files {slot[-slot][,....]}
```

Example

The following example shows how you view all debug files for all slots, and then remove the debug files for slot 1.

```
VSP-9012:1>show debug-file
```

```
=====
Core Files
=====
Directory: /intflash/coreFiles/1
1. File:   core.logServer.20120611084204.1.tar
   Size:   60928 bytes
   Created: Mon Jun 11 08:42:04 2012
2. File:   core.trcServer.20120611084213.1.tar
   Size:   60928 bytes
   Created: Mon Jun 11 08:42:13 2012
3. File:   core.logServer.20120611164647.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:46:48 2012
4. File:   core.trcServer.20120611164652.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:46:52 2012
5. File:   core.dbgServer.20120611164700.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:47:01 2012
6. File:   core.logServer.20120611164740.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:47:41 2012

Remote CP Directory: /intflash/coreFiles/2
1. File:   core.coreManager.x.20120612085548.2.tar
   Size:   1162240 bytes
   Created: Tue Jun 12 08:55:49 2012
2. File:   core.coreManager.x.20120612085602.2.tar
   Size:   478208 bytes
   Created: Tue Jun 12 08:56:02 2012
3. File:   core.coreManager.x.20120612085553.2.tar
   Size:   1170432 bytes
   Created: Tue Jun 12 08:55:56 2012
4. File:   core.coreManager.x.20120612085558.2.tar
   Size:   1883136 bytes
   Created: Tue Jun 12 08:56:00 2012

=====
Archive Files
=====
Directory: /intflash/archive/1
1. File:   archive.20120611083021.1.tar
```

Software troubleshooting

```
Size: 34296320 bytes
Created: Mon Jun 11 08:30:22 2012
2. File: archive.20120611163454.1.tar
Size: 31108096 bytes
Created: Mon Jun 11 16:34:54 2012
3. File: archive.20120611164354.1.tar
Size: 31792128 bytes
Created: Mon Jun 11 16:43:55 2012
4. File: archive.20120611164507.1.tar
Size: 31881216 bytes
Created: Mon Jun 11 16:45:08 2012

Remote CP Directory: /intflash//archive/2
1. File: archive.20120611163507.2.tar
Size: 30903296 bytes
Created: Mon Jun 11 16:35:08 2012
2. File: archive.20120611164408.2.tar
Size: 31314432 bytes
Created: Mon Jun 11 16:44:09 2012
3. File: archive.20120611164521.2.tar
Size: 31367168 bytes
Created: Mon Jun 11 16:45:21 2012

Directory: /intflash/archive/4
1. File: archive.20120611163515.4.tar
Size: 4725760 bytes
Created: Mon Jun 11 16:35:18 2012
2. File: archive.20120611164416.4.tar
Size: 5639168 bytes
Created: Mon Jun 11 16:44:20 2012
3. File: archive.20120611164529.4.tar
Size: 5760000 bytes
Created: Mon Jun 11 16:45:33 2012

Directory: /intflash/archive/SF4
1. File: archive.20120611163536.SF4.tar
Size: 1550336 bytes
Created: Mon Jun 11 16:35:40 2012
2. File: archive.20120611164436.SF4.tar
Size: 1781248 bytes
Created: Mon Jun 11 16:44:39 2012
3. File: archive.20120611164549.SF4.tar
Size: 1811968 bytes
Created: Mon Jun 11 16:45:53 2012
```

```
=====
PMEM Files
=====
```

```
Directory: /intflash/PMEM/4
1. File: pmem.20120607194023.4.bin.gz
Size: 571048 bytes
Created: Thu Jun 7 19:40:23 2012
```

```
=====
DMalloc Files
=====
```

```
=====
Flrec Files
=====
```

```
=====
WdStats Files
=====
```

```

Directory: /intflash/wd_stats/4
1. File:    wd_stats.log.backup
   Size:    2311 bytes
   Created: Mon Jun 11 09:25:07 2012

VSP-9012:1>debug-file remove 1
VSP-9012:1>show debug-file

=====
                               Core Files
=====
Remote CP Directory: /intflash/coreFiles/2
1. File:    core.coreManager.x.20120612085548.2.tar
   Size:    1162240 bytes
   Created: Tue Jun 12 08:55:49 2012
2. File:    core.coreManager.x.20120612085602.2.tar
   Size:    478208 bytes
   Created: Tue Jun 12 08:56:02 2012
3. File:    core.coreManager.x.20120612085553.2.tar
   Size:    1170432 bytes
   Created: Tue Jun 12 08:55:56 2012
4. File:    core.coreManager.x.20120612085558.2.tar
   Size:    1883136 bytes
   Created: Tue Jun 12 08:56:00 2012

=====
                               Archive Files
=====
Remote CP Directory: /intflash//archive/2
1. File:    archive.20120611163507.2.tar
   Size:    30903296 bytes
   Created: Mon Jun 11 16:35:08 2012
2. File:    archive.20120611164408.2.tar
   Size:    31314432 bytes
   Created: Mon Jun 11 16:44:09 2012
3. File:    archive.20120611164521.2.tar
   Size:    31367168 bytes
   Created: Mon Jun 11 16:45:21 2012

Directory: /intflash/archive/4
1. File:    archive.20120611163515.4.tar
   Size:    4725760 bytes
   Created: Mon Jun 11 16:35:18 2012
2. File:    archive.20120611164416.4.tar
   Size:    5639168 bytes
   Created: Mon Jun 11 16:44:20 2012
3. File:    archive.20120611164529.4.tar
   Size:    5760000 bytes
   Created: Mon Jun 11 16:45:33 2012

Directory: /intflash/archive/SF4
1. File:    archive.20120611163536.SF4.tar
   Size:    1550336 bytes
   Created: Mon Jun 11 16:35:40 2012
2. File:    archive.20120611164436.SF4.tar
   Size:    1781248 bytes
   Created: Mon Jun 11 16:44:39 2012
3. File:    archive.20120611164549.SF4.tar
   Size:    1811968 bytes
   Created: Mon Jun 11 16:45:53 2012

=====
                               PMEM Files
=====

```

```
=====
Directory: /intflash/PMEM/4
1. File:      pmem.20120607194023.4.bin.gz
   Size:      571048 bytes
   Created:   Thu Jun  7 19:40:23 2012

=====
                          DMalloc Files
=====

=====
                          Flrec Files
=====

=====
                          WdStats Files
=====
Directory: /intflash/wd_stats/4
1. File:      wd_stats.log.backup
   Size:      2311 bytes
   Created:   Mon Jun 11 09:25:07 2012
```

The following example shows how to view only core files on the switch.

```
VSP-9012:1#show core-files
=====
                          Core Files
=====
Directory: /intflash/coreFiles/1
1. File:      core.1353113115.lifecycle.CP.1.gz
   Size:      139406 bytes
   Created:   Fri Nov 16 19:45:15 2012
2. File:      core.cbcp-main.x.20121114043335.1.tar
   Size:      14059520 bytes
   Created:   Wed Nov 14 04:35:36 2012
3. File:      core.cbcp-main.x.20121114045202.1.tar
   Size:      12809728 bytes
   Created:   Wed Nov 14 04:54:03 2012
4. File:      core.cbcp-main.x.20121114050825.1.tar
   Size:      12638720 bytes
   Created:   Wed Nov 14 05:10:26 2012
5. File:      core.cbcp-main.x.20121114122506.1.tar
   Size:      13020160 bytes
   Created:   Wed Nov 14 12:27:07 2012
6. File:      core.1353336274.lifecycle.CP.1.gz
   Size:      139390 bytes
   Created:   Mon Nov 19 09:44:34 2012
7. File:      core.1353319337.lifecycle.CP.1.gz
   Size:      139404 bytes
   Created:   Mon Nov 19 05:02:17 2012
8. File:      core.cbcp-main.x.20130122182946.1.tar
   Size:      13683712 bytes
   Created:   Tue Jan 22 18:32:08 2013
9. File:      core.cbcp-main.x.20130220143809.1.tar
   Size:      13969920 bytes
   Created:   Wed Feb 20 14:38:10 2013
10. File:     core.cbcp-main.x.20130225155025.1.tar
   Size:      13526016 bytes
   Created:   Mon Feb 25 15:50:25 2013
11. File:     core.cbcp-main.x.20130225155407.1.tar
   Size:      12674560 bytes
   Created:   Mon Feb 25 15:54:07 2013
```


Variable definitions

Use the data in the following table to use the `show core-files` command.

Variable	Value
<code>{slot[-slot][,...]}</code>	Displays the core files for the slot that you select.

Use the data in the following table to use the `show debug-file` command.

Variable	Value
<code>all</code>	Displays all types of debug files.
<code>{slot[-slot][,...]}</code>	Displays debug files for the slot that you select. If you do not select a slot number, the device displays all types of the archived debug files present in a slot by slot basis. If you select a slot number, the device only displays archived files for the slot you select.

Use the data in the following table to use the `debug-file remove` command.

Variable	Value
<code>all</code>	Removes all types of debug files in all slots. If you use the option <code>all</code> with the <code>remove debug-file</code> command, then the device deletes all types of debug files, including the latest debug files.
<code>{slot[-slot][,...]}</code>	Removes debug files for the slot that you select. When you clear archived files, if you do not select a slot number, the device deletes all types of archived debug files except the latest file in each slot. Valid slots are 1–12, SF1–SF6, and all.

Configuring port mirroring

About this task

Use port mirroring to aid in diagnostic and security operations.

Use port mirroring to make a copy of a traffic flow and send that copy to a device for analysis, for example, for diagnostic sniffing. Use the mirror to see the packets in the flow without breaking into the physical connection to place a packet onto the sniffer inline. You can also use port mirroring for security. You can send flows to inspection engines for post processing.

Connect the sniffer (or other traffic analyzer) to the output port you specify in this procedure.

Configure a destination IP address to monitor for Layer 3 mirroring.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a port mirroring instance:

```
mirror-by-port <1-479> in-port {slot/port[-slot/port][,...]}
{monitor-ip {A.B.C.D} [dscp <0-63>] [ttl <2-255>]|monitor-mlt
<1-512>|monitor-vlan <1-4084>|out-port {slot/port[-slot/port]
[,...]]}
```

3. Configure the mode:

```
mirror-by-port <1-479> mode <both|rx|tx>
```

4. Enable the mirroring instance:

```
mirror-by-port <1-479> enable
```

5. Modify existing mirroring entries as required:

```
mirror-by-port mirror-port <1-479> {slot/port[-slot/port][,...]}
```

OR

```
mirror-by-port monitor-ip <1-479> {A.B.C.D} [dscp <0-63>] [ttl
<2-255>]
```

OR

```
mirror-by-port monitor-mlt <1-479> <1-512>
```

OR

```
mirror-by-port monitor-port <1-479> {slot/port[-slot/port][,...]}
```

OR

```
mirror-by-port monitor-vlan <1-479> <1-4084>
```

*** Note:**

Before you can modify an existing entry, you must disable the entry: `no mirror-by-port <1-479> enable`.

6. Verify the configuration:

```
show mirror-by-port [WORD<1-1024>]
```

Example

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

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

Create the port mirroring instance. Traffic passing port 5/15 mirrors to port 5/16:

```
VSP-9012:1(config)# mirror-by-port 8 in-port 5/15 out-port 5/16
```

The analyzer connects to port 5/16.

Disable the entry:

```
VSP-9012:1(config)# no mirror-by-port 8 enable
```

Mirror both ingress and egress traffic passing through port 5/16:

```
VSP-9012:1(config)# mirror-by-port 8 mode both
```

Enable mirroring for the instance:

```
VSP-9012:1(config)# mirror-by-port 8 enable
```

Configure Layer 3 mirroring:

```
VSP-9012:1(config)#mirror-by-port 8 in-port 5/16 monitor-ip 5.5.5.5 dscp
10 ttl 15
```

The following example shows sample command output; it does not necessarily reflect the preceding examples.

```
VSP-9012:1(config)#show mirror-by-port
```

```
=====
                          Diag Mirror-By-Port
=====
```

ID	MIRRORED_PORT	MIRRORING_DEST	ENABLE	MODE	REMOTE-MIRROR VLAN-ID	DSCP	TTL
1	5/1	5/2	true	rx	0	0	64
2	5/3	5/4	true	rx	0	0	64
3	5/5	5/6	true	rx	0	0	64
4	5/7	5/8	true	rx	0	0	64
5	5/9	5/10	true	rx	0	0	64
6	5/11	5/12	true	rx	0	0	64
7	5/13	5/14	true	rx	0	0	64
8	5/15	5/16	true	rx	0	0	64
9	5/17	5/18	true	rx	0	0	64
11	5/19	5/20	true	rx	0	0	64
12	5/21	5/22	true	rx	0	0	64
13	5/23	5/24	true	rx	0	0	64
14	5/25	5/26	true	rx	0	0	64
15	5/27	5/28	true	rx	0	0	64
16	5/29	5/30	true	rx	0	0	64
20	5/31	5/32	true	rx	0	0	64

```
16 out of 24 Total Num of MirIds displayed
```

```
VSP-9012:1(config)#show mirror-by-port 1,5,12-15,20
```

```
=====
                          Diag Mirror-By-Port
=====
```

ID	MIRRORED_PORT	MIRRORING_DEST	ENABLE	MODE	REMOTE-MIRROR VLAN-ID	DSCP	TTL
1	5/1	5/2	true	rx	0	0	64
5	5/9	5/10	true	rx	0	0	64
12	5/21	5/22	true	rx	0	0	64
13	5/23	5/24	true	rx	0	0	64
14	5/25	5/26	true	rx	0	0	64
15	5/27	5/28	true	rx	0	0	64
20	5/31	5/32	true	rx	0	0	64

```
7 out of 7 matched entries out of total 24 Mirror entries displayed.
```

Variable definitions

Use the data in the following table to use the **mirror-by-port** command.

Table 17: Variable definitions

Variable	Value
<1-479>	Specifies the entry ID.
enable	Enables or disables a mirroring instance already created in the mirror-by-port table.
in-port {slot/port[-slot/port][,...]} {monitor-ip {A.B.C.D} [dscp <0-63>] [ttl <2-255>] monitor-mlt <1-512> monitor-vlan <1-4084> out-port {slot/port[-slot/port][,...]}}	<p>Creates a new mirror-by-port table entry.</p> <ul style="list-style-type: none"> in-port {slot/port[-slot/port][,...]} specifies the mirrored port. monitor-ip {A.B.C.D} [dscp <0-63>] [ttl <2-255> specifies the destination IP address for Layer 3 remote mirroring. You can optionally configure the DSCP and time-to-live values, or accept the defaults. monitor-mlt <1-512> specifies the mirroring MLT ID from 1–512. monitor-vlan <1-4084> specifies the mirroring VLAN ID from 1–4084. out-port {slot/port[-slot/port][,...]} specifies the mirroring port.
mirror-port <1-479> {slot/port[-slot/port][,...]}	<p>Modifies the mirrored port.</p> <p>Before you can modify an existing entry, you must disable the entry: <code>no mirror-by-port <1-479> enable</code>.</p>
monitor-ip <1-479> {A.B.C.D} [dscp <0-63>] [ttl <2-255>]	<p>Creates a mirroring instance for Layer 3 remote mirroring. The destination must be an IP address {A.B.C.D}. The default DSCP is 0 and the default TTL is 255.</p> <p>For Layer 3 mirroring, every hop in the path from the mirrored port to the remote-mirroring port must be routed. Avaya recommends that you configure the remote-mirrored port in its own VLAN at the last hop to prevent flooding.</p>
monitor-mlt <1-479> <1-512>	<p>Modifies the monitoring MLT. <1-512> specifies the mirroring MLT ID.</p> <p>Before you can modify an existing entry, you must disable the entry: <code>no mirror-by-port <1-479> enable</code>.</p>
monitor-port <1-479> {slot/port[-slot/port][,...]}	<p>Modifies the monitoring ports.</p> <p>Before you can modify an existing entry, you must disable the entry: <code>no mirror-by-port <1-479> enable</code>.</p>
monitor-vlan <1-479> <1-4084>	Modifies the monitoring VLAN.

Table continues...

Variable	Value
	Before you can modify an existing entry, you must disable the entry: <code>no mirror-by-port <1-479> enable</code> .
<code>mode <both rx tx></code>	Configures the mirroring mode. The default is rx. <ul style="list-style-type: none"> • both mirrors both egress and ingress packets. • rx mirrors ingress packets. • tx mirrors egress packets.
<code>remote-mirror-vlan-id <1-4084></code>	Configures the remote mirror VLAN ID.

Use the data in the following table to use the `show mirror-by-port` command.

Variable	Value
<code>WORD<1-1024></code>	Displays mirror-by-port diagnostic information. Mirror ID list {1-479}.

Configuring global mirroring actions with an ACL

Before you begin

- The ACL exists.

About this task

Configure the global action to mirror packets that match an access control list (ACL).

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the global action for an ACL:

```
filter acl set <1-2048> global-action {monitor-dst-mlt <1-512>|
monitor-dst-ports {slot/port[-slot/port ][,...]}|monitor-dst-vlan
<1-4084>}
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
Configure the global action for an ACL:
VSP-9012:1(config)#filter acl set 200 global-action monitor-dst-mlt 20
```

Variable definitions

Use the data in the following table to use the `filter acl set` command.

Table 18: Variable definitions

Variable	Value
<1-2048>	Specifies an ACL ID from 1–2048.
default-action <deny permit>	Specifies the global action to take for packets that do not match an ACL.
global-action {monitor-dst-mlt PT_MLT<1–512> monitor-dst-ports {slot/port[-slot/port][,...]}} monitor-dst-vlan <1–4084>}	Specifies the global action to take for matching ACLs: <ul style="list-style-type: none"> • monitor destination MLT—Configures mirroring to a destination MultiLink Trunking (MLT) group. • monitor destination ports—Configures mirroring to a destination port or ports. • monitor destination VLAN—Configures mirroring to a destination VLAN.

Configuring ACE actions to mirror

Before you begin

- The access control entry (ACE) exists.

About this task

Configure actions to use filters for flow-based mirroring.

If you use the mirror action, ensure that you specify the mirroring destination: IP address, MLTs, ports, or VLANs.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure actions for an ACE:

```
filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-ip
{A.B.C.D} [dscp <0-63>] [ttl <2-255>]
```

OR

```
filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-
mlt <1-512>
```

OR

```
filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-
ports {slot/port[-slot/port ][,...]}
```

OR

```
filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-
vlan <1-4084>
```

3. Ensure the configuration is correct:

```
show filter acl action [<1-2048>] [<1-2000>]
```

Example

Configure the ACL to permit mirroring to destination IP 192.0.2.2, the DSCP to 10 and TTL to 155:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#filter acl ace action 901 1 permit monitor-dst-ip192.0.2.2 dscp 10 ttl
155
```

Variable definitions

Use the data in the following table to use the `filter acl ace action` command.

Table 19: Variable definitions

Variable	Value
<code>1-2048</code>	Specifies the ACL ID from 1–2048
<code>1-2000</code>	Specifies the ACE ID from 1–2000.
<code>monitor-dst-ip {A.B.C.D} [dscp <0–63>] [ttl <2–255>]</code>	Configures mirroring to a destination IP address for flow-based mirroring. The destination must be an IP address {A.B.C.D}. The default DSCP is 256 (disabled) and the default TTL is 64.
<code>monitor-dst-mlt <1–512></code>	Configures mirroring to a destination MLT group.
<code>monitor-dst-ports {slot/port[-slot/port][,...]}</code>	Configures mirroring to a destination port or ports.
<code>monitor-dst-vlan <1-4084></code>	Configures mirroring to a destination VLAN.
<code>{permit deny}</code>	Configures the action mode for security ACEs. The default value is permit.

Configuring Layer 2 remote mirroring

About this task

Use remote mirroring to monitor many ports from different switches using one network probe device.

To configure remote mirroring for Layer 3, see [Configuring port mirroring](#) on page 121.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Configure the mode for remote mirroring:

```
remote-mirroring mode <source|termination>
```

3. Configure the destination MAC for remote mirroring:

```
remote-mirroring dstMac <0x00:0x00:0x00:0x00:0x00:0x00> [ether-type
<0x00-0xffff>] [vlan-id <1-4084>]
```

4. Configure the source MAC for remote mirroring:

```
remote-mirroring srcMac <0x00:0x00:0x00:0x00:0x00:0x00> [ether-type
<0x00-0xffff>] [vlan-id <1-4094>]
```

5. Specify a port list for remote mirroring:

```
remote-mirroring port {slot/port[-slot/port ][,...]} [mode <source|
termination>] [dstMac <0x00:0x00:0x00:0x00:0x00:0x00>] [srcMac
<0x00:0x00:0x00:0x00:0x00:0x00>] [ether-type <0x00-0xffff>] [vlan-id
<1-4094>]
```

6. Enable remote mirroring:

```
remote-mirroring enable
```

7. Ensure that the remote mirroring configuration is correct:

```
show remote-mirroring interfaces gigabitEthernet [dstMac
<0x00:0x00:0x00:0x00:0x00:0x00>] [enable <false|true>] [ether-type
<0x00-0xffff>] [mode <source|termination>] [srcMac
<0x00:0x00:0x00:0x00:0x00:0x00>] [vlan-id <1-4084>]
```

Example

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

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

```
VSP-9012:1(config)#interface gigabitEthernet 3/16
```

Configure the mode for remote mirroring:

```
VSP-9012:1(config-if)#remote-mirroring mode source
```

Configure the destination MAC for remote mirroring:

```
VSP-9012:1(config-if)#remote-mirroring dstMac 00-C0-E0-86-AA-F7 ether-
type 07-00-2C vlan-id 20
```

Configure the source MAC for remote mirroring:

```
VSP-9012:1(config-if)#remote-mirroring srcMac 00-B0-E1-85-AA-E2 ether-
type 07-00-2C vlan-id 200
```

Specify a port list for remote mirroring:

```
VSP-9012:1(config-if)#remote-mirroring port 3/10-3/12
```

Enable remote mirroring:

```
VSP-9012:1(config-if)#remote-mirroring enable
```


Ensure that remote mirroring configuration is correct:

```
VSP-9012:1(config-if)#show remote-mirroring interfaces gigabitEthernet
```

```
=====
Port Remote Mirroring
=====
PORT   Enable MODE      SourceMac          DestinationMac     EtherType Vid-List
-----
3/16   source           00-B0-E1-85-AA-E2  00-C0-E0-86-AA-F7  07-00-2C
200
```

Variable definitions

Use the data in the following table to use the `remote-mirroring` and `show remote-mirroring interfaces` commands.

Table 20: Variable definitions

Variable	Value
<code>dstMac <0x00:0x00:0x00:0x00:0x00:0x00></code>	Configures the destination MAC address for use in the remote mirroring encapsulation header. The device sends the mirrored packet to this MAC address. Only remote mirroring source (RMS) ports use the DstMac. For remote mirroring termination (RMT) ports, the device uses one of the unused MAC addresses from the switch port MAC address range. The device saves this MAC address in the configuration file.
<code>enable</code>	Enables remote mirroring on the port. When remote mirroring is enabled, the following events occur: <ul style="list-style-type: none"> • A device adds a static entry for the DstMac to the forwarding database. The switch sends all packets that use this remote mirroring DstMac to the RMT port. • The switch periodically (once in 10 seconds) transmits broadcast Layer 2 packets in the associated VLAN so that all nodes in the network can learn the DstMac.
<code>ether-type <0x00-0xffff></code>	Specifies the Ethertype of the remote mirrored packet. The default value is 0x8103.
<code>mode <source termination></code>	Specifies whether the port is an RMT (mode is termination) or an RMS (mode is source).
<code>{slot/port[-slot/port][,...]}</code>	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
<code>srcMac <0x00:0x00:0x00:0x00:0x00:0x00></code>	Configures the source MAC address for use in the remote mirroring encapsulation header. The device

Table continues...

Variable	Value
	sends the mirrored packet from the RMS port, and the device derives the source MAC parameter in the header from this address. The source MAC address of the encapsulated frame contains the first 45 bits of this MAC address. The device derives the three least significant bits from the port number of the RMS port. The default value is the MAC address of the port.
vlan-id <1-4084>	Specifies to which VLAN the remote mirroring destination MAC address belongs. The VLAN must be a port-based VLAN. Use this variable only for RMT ports. After you remove the RMT port from the last VLAN in the list, RMT is disabled on the port.

Accessing the secondary CPU

Before you begin

- The secondary CPU has an IP address.

About this task

Access the secondary CPU to gain access to the PCAP engine. You can gain access to the PCAP engine through a direct console connection to the secondary CPU, or by using a peer Telnet session from the master CPU.

Procedure

1. Log on to the primary CPU.
2. Enter Privileged EXEC mode:
enable
3. Access the secondary CPU:
peer telnet

Example

Access the secondary CPU:

```
VSP-9012:1>enable
VSP-9012:1#peer telnet
```

Configuring PCAP global parameters

Configure PCAP globally to define how PCAP operates on the Avaya Virtual Services Platform 9000.

*** Note:**

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Before you begin

- The secondary CP module is active.
- If you save to external flash, a Compact Flash (CF) card is present.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable PCAP:

```
pcap enable
```

3. Configure optional parameters as required.

4. Ensure the configuration is correct:

```
show pcap
```

Example

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

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

Enable PCAP globally:

```
VSP-9012:1(config)# pcap enable
```

Enable buffer wrapping:

```
VSP-9012:1(config)# pcap buffer-wrap
```

Enable external flash wrapping:

```
VSP-9012:1(config)# pcap wrap-auto-save-file
```

Specify the buffer size to 32 MB:

```
VSP-9012:1(config)# pcap buffer-size 32
```

Specify the fragment-size to 64 bytes:

```
VSP-9012:1(config)# pcap fragment-size 64
```

Enable auto-save to the external flash with the file name pcap.cap:

```
VSP-9012:1(config)# pcap auto-save file-name pcap.cap extflash
```

Display PCAP settings:

```
VSP-9012:1(config)# show pcap
```

```
VSP-9012:1(config)#show pcap
  enable = TRUE
  buffer-wrap = TRUE
  wrap-auto-save-file = TRUE
  buffer-size = 32 MB
  fragment-size = 64 Bytes
  auto-save = TRUE
  AutoSaveFilename = pcap.cap
  AutoSaveDevice = extflash
```

Variable definitions

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

Table 21: Variable definitions

Variable	Value
auto-save [file-name <i>WORD</i> <1-40>] [network { <i>A.B.C.D</i> }extflash]	Enables or disables auto-save. If you enable auto-save, PCAP saves the captured frames into the device you specify and continues to capture frames. The default is enable. If you disable this option, PCAP stores packets in the DRAM buffer only. <ul style="list-style-type: none"> file-name <i>WORD</i><1-40> is the name of the file where the device saves the captured frames. network configures the save device to network. {<i>A.B.C.D</i>} is the IP address of the network device. extflash configures the save device to a Compact Flash card.
buffer-size <2-128>	Specifies the size of the buffer for storing data. The default is 32 MB.
buffer-wrap	Enables buffer wrapping. When this variable is true and the buffer becomes full, the capture continues by wrapping the buffer. If this variable is false and the buffer becomes full, the packet capture stops. The default value is true. The device generates a log message after the buffer wraps.
enable	Enables PCAP globally. The default is disabled. To disable PCAP, use the <code>no pcap enable</code> command.
fragment-size <64-9600>	Specifies the number of bytes from each frame to capture. The default is the first 64 bytes of each frame.
reset-stat	Resets the PCAP engine DRAM buffer, as well as all software counters used for PCAP statistics. You can execute this command in the primary or secondary CPU.
wrap-auto-save-file	Enables wrap around auto-save-file. When this variable is true and the autosave device is extflash, this causes an overwrite of the present file on the Compact Flash card during an autosave. The system generates a log after the file is overwritten on the Compact Flash card.

Table continues...

Variable	Value
	If this variable is false, the present file is not overwritten on the Compact Flash card.

Enabling PCAP on a port

Configure PCAP on a port to capture packets on that port.

* Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Before you begin

- If required, IP filters exist.
- If required, ACLs with a global action of mirror exist.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[-slot/port] [, ...]} OR interface
vlan <1-4084>
```

2. Enable PCAP on ports:

```
pcap enable [mode {both|rx|tx}]
```

3. Verify the PCAP configuration:

```
show pcap port
```

Example

Enable PCAP on port 3/5 for ingress and egress packets:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitEthernet 3/5
VSP-9012:1(config-if)#pcap enable mode both
```

Variable definitions

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

Table 22: Variable definitions

Variable	Value
enable [mode {both rx tx}]	Enables or disables PCAP on the port. The default PCAP mode captures ingress packets (rx mode). If you enable PCAP in filter mode, then the device only captures packets that match the filter criteria.

Configuring PCAP capture filters

Use capture filters to better define the match criteria used on packets.

Avaya recommends that you use PCAP with IP or MAC filters to reduce the load on the PCAP engine.

To create a functional capture filter that captures specific packets, create two filters. Use one filter to capture specific packets, and another filter to drop all other packets.

* Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a capture filter:

```
pcap capture-filter <1-1000>
```

3. Configure the filter action:

```
pcap capture-filter <1-1000> action <capture|drop|trigger-off|
trigger-on>
```

4. Define the match parameters.

Use the following variable definitions table to configure match parameters.

5. Enable the filter:

```
pcap capture-filter <1-1000> enable
```

6. Ensure the configuration is correct:

```
show pcap capture-filter [<1-1000>]
```

Example

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

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

Configure the filter action to capture:

```
VSP-9012:1(config)# pcap capture-filter 2
```

```
VSP-9012:1(config)# pcap capture-filter 2 action capture
```

Specify the DSCP value in a range of 1 to 63:

```
VSP-9012:1(config)# pcap capture-filter dscp 1 63
```

Specify the destination IP address:

```
VSP-9012:1(config)# pcap capture-filter dstip 46.12.17.11
```

Specify the source IP address:

```
VSP-9012:1(config)# pcap capture-filter srcip 45.10.17.10
```

```
VSP-9012:1(config)# pcap capture-filter 2 enable
```

Variable definitions

Use the data in the following table to use the `pcap capture-filter` command.

Table 23: Variable definitions

Variable	Value
<1-1000>	Specifies the capture filter ID.
action <capture drop trigger-off trigger-on>	<p>Determines the action taken by the filter:</p> <ul style="list-style-type: none"> capture indicates that the packet is captured. drop indicates that the packet is dropped. trigger-off indicates that PCAP captures packets until one matches the criteria, and then disables the filter entr, and globally disables PCAP. trigger-on indicates that PCAP captures a packet after it matches the criteria, and then disables the filter entry. <p>The default is capture.</p> <p>! Important:</p> <p>Because the PCAP engine runs on the secondary CP module, the master CP module does not reflect the change in PCAP and PCAP capture-filter status if you use the action trigger-on or trigger-off. Run the <code>show pcap capture-filter</code> or <code>show pcap cli</code> commands on the secondary CP module to view the correct status. After PCAP disables</p>

Table continues...

Variable	Value
	the filter entry on the PCAP engine, if you use the <code>show pcap</code> or <code>show pcap capture-filter</code> command on the master CP module, the status appears as true (enabled), when it really is false (disabled). To activate PCAP and the PCAP capture filter again, you must reenable them on the master CP module.
dscp <0-63> [<0-63>] [match-zero]	Specifies the DSCP value of the packet. <0-63> is the DSCP from 0–63. The default is 0, which means this option is disabled. Use the second <0-63> to specify a range. When you configure match-zero, 0 is considered a valid value; otherwise, 0 is considered a disable value.
dstip < <A.B.C.D>[<A.B.C.D>]	Specifies the destination IP address. The default is 0.0.0.0, which means this option is disabled. Use the second <A.B.C.D> to specify a range.
dstmac <0x00:0x00:0x00:0x00:0x00:0x00> [<1-6>]	Specifies the MAC address of the destination. If you configure the mask, then the device only compares the first few bytes. <1-6> is the destination MAC address mask, and specifies a range.
enable	Enables the filter. The default is disable.
ether-type <0x0-0xffff> [<0x0-0xffff>]	Specifies the Ethertype of the packet. The default is 0, meaning that this option is disabled. Use the second <0x0-0xffff> to specify a range.
packet-count <0-65535>	Stops PCAP after capturing the specified number of packets. This variable is similar to the refresh-timer option; after it is invoked, the filter is disabled. This option is active only if you configure the action parameter to trigger-on. The default value is 0, which means this option is disabled.
pbits <0-7> [<0-7>]	Specifies the priority bit of the packet. The default is 0, which means this option is disabled. Use the second <0-7> to specify a range. When match-zero is set, 0 is considered a valid value; otherwise, 0 is considered a disable value.
protocol-type<0-255><0-7>[<0-255>]	Specifies the packet protocol type. The default is 0, which means this option is disabled. Use the second <0-255> to specify a range.
refresh-timer WORD<1-7>	Starts or restarts the timer. After the PCAP engine receives a matching packet, it disables the capture filter. If the PCAP engine

Table continues...

Variable	Value
	does not receive another matching packet within the specified time, PCAP is disabled globally. The timer restarts every time the PCAP engine receives a packet, until the timer expires. Specify the value in milliseconds (ms). This variable is active only if the filter action is trigger-on. To delete this option, configure it to 0. The default value is 0.
srcip <A.B.C.D>[<A.B.C.D>]	Specifies the source IP address. The default is 0.0.0.0, which means this option is disabled. Use the second <A.B.C.D> to specify a range.
srcmac<0x00:0x00:0x00:0x00:0x00:0x00>[<1-6>]	Specifies the MAC address of the source. If you configure the mask, then the device only compares the first few bytes. The default is 00:00:00:00:00:00, which means this option is disabled. <1-6> is the mask of the source MAC address. This option specifies an address range.
tcp-port<0-65535> [<0-65535>]	Specifies the TCP port of the packet. The default is 0, which means this option is disabled. Use the second <0-65535> to specify a range.
timer <i>WORD</i> <1-7>	Specifies that PCAP is invoked after the first packet matches and stops after a configured value of time. After the timer starts, the filter is disabled. After the PCAP engine receives a matching packet, it captures all packets for the duration of the timer, and then disables PCAP globally. This option is active only if the filter action is trigger-on. <i>WORD</i> <1-7> is a value from 100-3600000 milliseconds. The default value is zero. Configure the value to 0 to disable the timer.
udp-port <0-65535> [<0-65535>]	Specifies the UDP port of the packet. The default is 0, which means this option is disabled. Use the second <0-65535> to specify a range.
user-defined <0-9600> <i>WORD</i> <1-50>	Configures a user-defined value on which to match the packet. You can define a pattern in hexadecimal or characters to match (<0-9600>). You can also specify the offset to start the match (<i>WORD</i> <1-50>). The default value of pattern is null ("), which means that this field is discarded. To disable this option, configure the pattern to null (").
vid<1-4084>[<1-4084>]	Specifies the VLAN ID of the packet. The default is 0, which means that this option is disabled. Use the second <1-4084> to specify a range.

Using the captured packet dump

You can view packets using an ACLI session and the secondary CPU. Dumping a large number of captured packets is CPU intensive. The device does not respond to commands while the dump is in progress. Avaya recommends that you use this command only when absolutely necessary. However, no degradation in normal traffic handling or switch failover occurs.

* Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Procedure

1. Log on to the secondary CPU.
2. Use the following command:

```
show pcap dump
```

Copying captured packets to a remote machine

About this task

You can copy packets to a remote machine, or an internal, or an external flash. If you use PCAP with autosave disabled, the device stores the captured packets in the secondary CPU dynamic random-access memory (DRAM) buffer. To copy the packets to a file for later viewing, use the `copy` command. You can use this command in the primary CPU.

Captured packets stored in the secondary CPU DRAM buffer use the name PCAP00.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. Copy packets:

```
copy WORD<1-99> WORD<1-99>
```
3. Copy packets from DRAM:

```
copy PCAP00 WORD<1-99>
```
4. Use File Transfer Protocol (FTP) to transfer the file for later viewing:

```
ftp>get PCAP00 WORD<1-99>
```

Example

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

```
VSP-9012:1# copy 46.11.10.33/pcap.cap /intflash/pcap.cap
```

```
VSP-9012:1# copy PCAP00 /inflash/pcap.cap
```

Use File Transfer Protocol (FTP) to transfer the file for later viewing:

```
ftp> get PCAP00 45.16.11.34
```

Variable definitions

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

Table 24: Variable definitions

Variable	Value
<code>WORD<1-99> WORD<1-99></code>	<p>Specifies USB, flash, or an IP host by IP address and specifies the PCAP file (.cap). Formats include:</p> <ul style="list-style-type: none"> • a.b.c.d:<file> • /intflash/<file> • /mnt/intflash/<file> • /extflash/<file> • /mnt/extflash/<file> • /usb/<file> <p>The first <code>WORD<1-255></code> identifies the source location and file name. The second <code>WORD<1-255></code> identifies the destination location and file name.</p> <p>/mnt/intflash is equivalent to the /intflash of the standby CP (if present).</p> <p>/mnt/extflash is equivalent to the /extflash of the standby CP (if present).</p>

Resetting the PCAP DRAM buffer

Reset the DRAM buffer to clear the PCAP dynamic random-access memory (DRAM) buffer and the PCAP counters.

Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Procedure

1. Log on to the secondary CPU.
2. Enter Global Configuration mode:

- ```
enable
```
- ```
configure terminal
```
3. Disable PCAP:


```
no pcap enable
```
 4. Reset the PCAP engine DRAM buffer:


```
pcap reset-stat
```
 5. Reenable PCAP:


```
pcap enable
```

Clearing ARP information for an interface

About this task

Clear the Address Resolution Protocol (ARP) cache as part of ARP problem resolution procedures.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Clear ARP information:


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

OR

```
clear ip arp interface vlan <1-4084>
```

Example

Clear ARP information:

```
VSP-9012:1>enable
VSP-9012:1#clear ip arp interface gigabitethernet 4/12
```

Variable definitions

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

Table 25: Variable definitions

Variable	Value
<code>1-4084</code>	Specifies the VLAN ID.
<code>slot/port[-slot/port][,...]</code>	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Flushing routing, MAC, and ARP tables for an interface

About this task

Flush or clear the routing tables for administrative and troubleshooting purposes. The clear and flush commands perform the same function; they remove the contents of the table.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Flush IP routing tables by port:

```
action flushIp
```

3. Flush the MAC address tables:

```
action flushMacFdb
```

4. Flush ARP tables:

```
action flushArp
```

5. Flush all tables with one command:

```
action flushAll
```

6. Exit to Global Configuration mode:

```
exit
```

7. Clear a routing table:

```
clear ip route gigabitethernet {slot/port}
```

OR

```
clear ip route vlan <1-4084>
```

Example

Flush all tables:

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

Variable definitions

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

Table 26: Variable definitions

Variable	Value
1–4084	Specifies the VLAN ID.
{slot/port}	Specifies a port number.

Pinging an IP device

Ping a device to test the connection between the Avaya Virtual Services Platform 9000 and another network device.

About this task

After you ping a device, the switch sends an Internet Control Message Protocol (ICMP) packet to the target device. If the device receives the packet, it sends a ping reply. After the switch receives the reply, a message appears that indicates traffic can reach the specified IP address. If the switch does not receive a reply, the message indicates the address does not respond.

Ping and traceroute can fail for VRF routes if you use large packet sizes for the operation. Do not use packet sizes larger than the following:

- Ping for VRF Lite: 1480 bytes
- Traceroute for VRF Lite: 1444 bytes

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Ping an IP network connection:

```
ping WORD<0-256> [-d] [-I <1-60>] [-s] [-t <1-120>] [count <1-9999>]
[datasize {28-9216|28-51200}] [interface WORD<1-256>|
gigabitEthernet|mgmtEthernet|tunnel|vlan] [scopeid <1-9999>] [source
WORD<1-256>] [vrf WORD<1-16>]
```

Example

Ping an IP device through the management interface.

```
VSP-9012:1>ping 46.16.10.35 vrf mgmtrouter
46.16.10.35 is alive
```

Variable definitions

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

Variable	Value
count <1-9999>	Specifies the number of times to ping (1–9999).
-d	Configures the ping debug mode. This variable detects local software failures (ping related threads creation or write to sending socket) and receiving

Table continues...

Variable	Value
	issues (ICMP packet too short or wrong ICMP packet type).
datasize {28-9216 28-51200}	Specifies the size of ping data sent in bytes. The datasize for IPv4 addresses is <28-9216>. The datasize for IPv6 addresses is <28-51200>. The default is 0.
-l <1-60>	Specifies the interval between transmissions in seconds (1-60).
interface WORD<1-256>{gigabitEthernet mgmtEthernet tunnel vlan}	Specifies the IP address of the outgoing interface. Additional ping interface parameters: <ul style="list-style-type: none"> • gigabitEthernet: {slot/port} gigabitethernet port • mgmtEthernet: {slot/port} management ethernet port • tunnel: tunnel ID as a value from 1 to 2147277248 • vlan: VLAN ID as a value from 1 to 4094
-s	Configures the continuous ping at the interval rate defined by the [-l] parameter.
scopeid <1-9999>	Specifies the circuit ID for IPv6.
source WORD<1-256>	Specifies the source IP address for the ping command.
-t <1-120>	Specifies the no-answer timeout value in seconds (1-120).
WORD<0-256>	Specifies the host name or IPv4 (a.b.c.d) or IPv6 (x:x:x:x:x:x) address (string length 0-256).
vrf WORD<1-16>	Specifies the virtual router and forwarder (VRF) name from 1-16 characters. Specify the MgmtRouter VRF if you need to run the ping operation through the management interface.

Running a traceroute test

Use traceroute to determine the route packets take through a network to a destination.

About this task

Ping and traceroute can fail for VRF routes if you use large packet sizes for the operation. Do not use packet sizes larger than the following:

- Ping for VRF Lite: 1480 bytes
- Traceroute for VRF Lite: 1444 bytes

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Run a traceroute test:

```
traceroute WORD<0-256> [<1-1176>] [-m <1-255>] [-p <1-65535>] [-q  
<1-255>] [-v] [-w <1-255>] [source WORD<1-256>] [vrf WORD<1-16>]
```

To use the source and vrf parameters in the same command, specify the VRF name before you specify the source address.

Example

Run traceroute test, with a probe packet size of 200 and a maximum time to live of 60.

```
VSP-9012:1>enable  
VSP-9012:1#traceroute 46.11.10.33 200 -m 60
```

Run a traceroute test for IPv6 address 2001:db8::.

```
VSP-9012:1>enable  
VSP-9012:1#traceroute 2001:db8::
```

Variable definitions

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

Table 27: Variable definitions

Variable	Value
-m <1-255>	Specifies the maximum time-to-live (TTL). The default is 30.
-p <1-65535>	Specifies the base UDP port number.
-q <1-255>	Specifies the number of probes for each TTL.
-v	Specifies verbose mode (detailed output).
-w <1-255>	Specifies the wait time for each probe.
<1-1176>	Specifies the size of the probe packet.
source <WORD 1-256>	Specifies the source IP address. If you do not specify a source address, the traceroute uses the primary IP address for the interface that sends the probe packet.
WORD<0-256>	Specifies the destination IPv4 or IPv6 address.
vrf <WORD 1-16>	Specifies the VRF instance by VRF name. This parameter applies only to IPv4.

Showing SNMP logs

About this task

Show the full SNMP logs. SNMP logs display the alarms and events registered on the device.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```
2. Show the logs:

```
show fulltech file WORD<1-99>
```

Variable definitions

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

Table 28: Variable definitions

Variable	Value
file <i>WORD</i> <1-99>	This variable represents the log file to open and display. Use one of three formats: <ul style="list-style-type: none"> • /intflash/<file> • /extflash/<file> • /usb/<file>

Using trace to examine IS-IS control packets

Use trace as a debug tool to examine the code flow and Intermediate-System-to-Intermediate-System (IS-IS) control packets. When you enable IS-IS trace flags, only trace information about the set flag appears.

Before you begin

 **Caution:**

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

About this task

Use the `trace level 119 <0-4>` command to trace IS-IS module information, including ACLI, instrumentation, High Availability, show config and platform dependent code. The IS-IS module ID is 119.

Use the `trace spbm isis level` command to trace platform independent code, IS-IS protocol, IS-IS hello, IS-IS adjacency, LSP processing, and IS-IS computation.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable the Intermediate-System-to-Intermediate-System trace flags:

```
trace flags isis set { none | tx-hello | rx-hello | tx-pkt | rx-pkt  
| adj | opt | tx-lsack | rx-lsack | tx-lsp | rx-lsp | pkt-err | nbr-  
mismatch | flood | spf-intra | spf-inter | spf-extern | prefix |  
nbr-change | intf-change | decide | fdb | dr | dd-masterslave |  
auth-fail | config | purge | policy | redist | tx-snp | rx-snp |  
timer | spbm-decide | global | perf | ha | ucast-fib | node | mcast-  
fib | isid | ip-shortcut|debug|ip-multicast}
```

3. Identify the module ID for which you want to use the trace tool:

```
show trace modid-list
```

4. Clear the trace:

```
clear trace
```

5. Turn on trace for IS-IS common code:

```
trace spbm isis level <0-4>
```

Wait approximately 30 seconds, and then stop trace.

6. Stop tracing:

```
trace shutdown
```

7. Display the trace information for SPBM IS-IS:

```
show trace spbm isis
```

8. Turn on trace for IS-IS in platform code:

```
trace level 119 [<0-4>]
```

Wait approximately 30 seconds.

The default trace settings for CPU utilization are:

- High CPU Utilization: 90%
- High Track Duration: 5 seconds
- Low CPU Utilization: 75%
- Low Track Duration: 5 seconds

 **Note:**

Module ID 119 represents the IS-IS module.

9. View the trace results:

```
trace screen enable
```

! Important:

If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

10. Save the trace file to the Compact Flash card for retrieval:

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

11. Display trace lines saved to a file:

```
show trace file [tail]
```

12. Search trace results for a specific string value:

```
trace grep [WORD<0-128>]
```

If you use this command and do not specify a string value, you clear the results of a previous search.

13. Stop tracing:

```
trace shutdown
```

14. Disable the Intermediate-System-to-Intermediate-System trace flags:

```
trace flags isis remove { none | tx-hello | rx-hello | tx-pkt | rx-
pkt | adj | opt | tx-lsack | rx-lsack | tx-lsp | rx-lsp | pkt-err |
nbr-mismatch | flood | spf-intra | spf-inter | spf-extern | prefix |
nbr-change | intf-change | decide | fdb | dr | dd-masterslave |
auth-fail | config | purge | policy | redist | tx-snp | rx-snp |
timer | spbm-decide | global | perf | ha | ucast-fib | node | mcast-
fib | isid | ip-shortcut }
```

Example

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

Clear prior trace information:

```
VSP-9012:1#clear trace
```

Enable IS-IS trace flags for received IS-IS hello packets:

```
VSP-9012:1#trace flags isis set rx-hello
```

Enable IS-IS trace flags for transmitted IS-IS hello packets:

```
VSP-9012:1#trace flags isis set tx-hello
```

Configure the module ID to 119 (IS-IS module) and the trace to 4 (very verbose):

```
VSP-9012:1#trace level 119 4
```

Enable the display of trace output to the screen:

```
VSP-9012:1#trace screen enable
```

```
VSP-9012:1# Screen tracing is on
```

Disable the display of trace output to the screen:

```
VSP-9012:1#trace screen disable
```

```
VSP-9012:1# Screen tracing is off
```

Variable definitions

Use the data in the following table to use the `trace flags isis` command.

Variable	Value
remove { none tx-hello rx-hello tx-pkt rx-pkt adj opt tx-lsack rx-lsack tx-lsp rx-lsp pkt-err nbr-mismatch flood spf-intra spf-inter spf-extern prefix nbr-change intf-change decide fdb dr dd-masterslave auth-fail config purge policy redist tx-snp rx-snp timer spbm-decide global perf ha ucast-fib node mcast-fib isid ip-shortcut debug ip-multicast }	Removes the Intermediate-System-to-Intermediate-System (IS-IS) trace flags for the specified option.
set { none tx-hello rx-hello tx-pkt rx-pkt adj opt tx-lsack rx-lsack tx-lsp rx-lsp pkt-err nbr-mismatch flood spf-intra spf-inter spf-extern prefix nbr-change intf-change decide fdb dr dd-masterslave auth-fail config purge policy redist tx-snp rx-snp timer spbm-decide global perf ha ucast-fib node mcast-fib isid ip-shortcut debug ip-multicast }	<p>Enables the Intermediate-System-to-Intermediate-System (IS-IS) trace flags for the specified option.</p> <ul style="list-style-type: none"> • none — none • tx-hello — Transmitted IS-IS hello packet • rx-hello — Received IS-IS hello packet • tx-pkt — Transmitted packet • rx-pkt — Received packet • adj — Adjacencies • opt — IS-IS TLVs • tx-lsack — Transmitted LSP acknowledgement • rx-lsack — Received LSP acknowledgement • tx-lsp — Transmitted Link State Packet • rx-lsp — Received Link State Packet • pkt-err — Packet Error • nbr-mismatch — Neighbor mismatch • flood — Flood • spf-intra — <i>Not used</i> • spf-inter — Shortest Path First Internal • spf-extern — Shortest Path First External


Table continues...

Variable	Value
	<ul style="list-style-type: none"> • prefix — Prefix • nbr-change — Neighbor change • intf-change — IS-IS circuit (interface) events • decide — Shortest Path First computation • fdb — Filtering Database • dr — Designated Router • <i>dd-masterslave</i> — <i>Not used</i> • auth-fail — Authorization failed • config — Configuration • purge — Link State Packet purge • <i>policy</i> — <i>Not used</i> • redistrib — Redistribute • tx-snp — Transmitted Sequence Number PDU (CSNP and PSNP) • rx-snp — Received Sequence Number Packet (CSNP and PSNP) • timer — Timer • spbm-decide — Shortest Path First computation for SPBM • <i>global</i> — <i>Not used</i> • perf — SPBM performance • ha — High Availability • ucast-fib — Unicast Forwarding Information Base • node — Node • mcast-fib — Multicast Forwarding Information Base • isid — I-SID • ip-shortcut — IP Shortcut • debug — Debug • ip-multicast — IP multicast

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


Variable	Value
file [tail]	Displays the trace results saved to a file.
level	Displays the current trace level for all modules.
modid-list	Specifies the module ID list.

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

Variable	Value
grep [<i>WORD</i> <0-128>]	Specifies the search keyword. You can use a specific MAC address. You can search for errors, using the command, <code>trace grep error</code> .
level [<0-192>] [<0-4>]	Starts the trace by specifying the module ID and level. <0-192> specifies the module ID. Module ID 119 represents the IS-IS module. Specifies the trace level: <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable enable}	Enables or disables the display of trace output to the screen.  Important: Avaya recommends you avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

Use the data in the following table to use the `save trace` command.

Variable	Value
file <i>WORD</i> <1-99>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d: <file> • /intflash/ <file> • /extflash/ <file> • /mnt/intflash/ <file> • /mnt/extflash/ <file> • /usb/<file> /mnt/intflash/ is the internal flash of the second CP module (the one to which you are not connected.) /mnt/extflash/ is the external flash of the second CP module (the one to which you are not connected.) WORD<1-99> is a string of 1-99 characters.

Variable	Value
	<p> Note:</p> <p>If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.</p>

Software troubleshooting tool configuration using EDM

Use the tools described in this section to perform troubleshooting procedures using Enterprise Device Manager (EDM).

Flushing routing tables by VLAN

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. Perform this procedure to flush the IP routing table for a VLAN.

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** box for the VLAN you want to flush, double-click, and then select a flush option from the list.

In a VLAN context, the device flushes all entries associated with the VLAN. You can also flush the Address Resolution Protocol (ARP) entries and IP routes for the VLAN.

5. Click **Apply**.

Flushing routing tables by port

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. Perform this procedure to flush the IP routing table for a port.

Procedure

1. Select a port.
2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the Interface tab.

- In the **Action** section, select **flushAll**.

In a port context, the device flushes all entries associated with the port. You can flush the ARP entries and IP routes for a port. After you flush a routing table, the routing table is not automatically repopulated. The repopulation time delay depends on the routing protocols in use.

- Click **Apply**.

Configuring port mirroring

Before you begin

- To change a port mirroring configuration, first disable mirroring.

About this task

Use port mirroring to aid in diagnostic and security operations.

Use port mirroring to make a copy of a traffic flow and send that copy to a device for analysis, for example, for diagnostic sniffing. Use the mirror to see the packets in the flow without breaking into the physical connection to place a packet onto the sniffer inline. You can also use port mirroring for security. You can send flows to inspection engines for post processing.

Connect the sniffer (or other traffic analyzer) to the output port you specify in this procedure.

Configure a destination IP address to configure Layer 3 remote mirroring.

Procedure

- In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
- Click **General**.
- Click the **Port Mirrors** tab.
- Click **Insert**.
- Configure mirroring as required.
- To enable port mirroring for the instance, select the **Enable** check box.
- Click **Insert**.

Port Mirrors field descriptions

Use the data in the following table to use the **Port Mirrors** tab.

Name	Description
Id	Specifies an assigned identifier for the configured port mirroring instance.
MirroredPortList	Specifies a port to be mirrored (the source port).

Table continues...

Name	Description
Enable	Enables or disables this port mirroring instance. The default value is Enable.
Mode	Specifies the traffic direction of the packet being mirrored: <ul style="list-style-type: none"> • tx mirrors egress packets. • rx mirrors ingress packets. • both mirrors both egress and ingress packets. The default is rx.
MirroringPortList	Specifies a destination port (the port to which the mirrored packets forward). Configures the mirroring port.
MirroringVlanId	Specifies the destination VLAN ID.
MirroringMltId	Specifies the destination MultiLink Trunking ID.
RemoteMirrorVlanId	Specifies the virtual local area network (VLAN) ID to which mirrored packets must be sent for remote mirroring. If set, this VLAN ID is used in the mirror tag of the remote mirrored packet.
MirroringIpAddr	Specifies the destination IP address for Layer 3 remote mirroring. For Layer 3 mirroring, every hop in the path from the mirrored port to the remote-mirroring port must be routed. Avaya recommends that you configure the remote-mirrored port in its own VLAN at the last hop to prevent flooding.
MirroringIpTtl	Specifies, optionally, the time-to-live value. The default is 64.
MirroringIpDscp	Specifies, optionally, the DSCP value. The default is 0.

Configuring Layer 2 remote mirroring

About this task

Use Layer 2 remote mirroring to monitor many ports from different switches using one network probe device.

Every hop in the path from the mirrored port to the remote-mirroring port must be routed. Avaya recommends that you configure the remote-mirrored port in its own VLAN at the last hop to prevent flooding.

To configure Layer 3 remote mirroring, see [Configuring port mirroring](#) on page 152.

Procedure

1. From the Physical Device View, select a port.
2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **Remote Mirroring** tab.
5. To add an entry, click **Insert**.

6. Select **Enable**.
7. Choose the mode.
8. Type the source MAC address (optional).
9. Type the destination MAC address.
10. Select a VLAN from the list (optional).
11. Click **Insert**.

Remote Mirroring field descriptions

Use the data in the following table to use the **Remote Mirroring** tab.

Name	Description
Index	Specifies the port.
Enable	<p>Enables or disables remote mirroring on the port. When remote mirroring termination (RMT) is enabled, the following things occur:</p> <ul style="list-style-type: none"> • The device adds a static entry for the DstMac to the FDB. The device sends all packets that come with that remote mirroring dstmac to the RMT port. • The switch periodically (once in 10 seconds) transmits broadcast Layer 2 packets in all associated VLANs so that all nodes in the network can learn the DstMac address. <p>The default is disabled.</p>
Mode	Specifies whether the port is a remote mirroring termination (RMT) or a remote mirroring source (RMS). The default is source.
SrcMac	Specifies the source MAC address of the remote mirrored packet. The device sends the remote mirroring packet with this source MAC address.
DstMac	Specifies the destination MAC address of the remote mirrored packet. The device bridges packets to this MAC address. The device sends remote mirroring packets to this MAC address.
EtherType	Specifies the EtherType of the remote mirrored packet. The device sends packets with this EtherType. The default value is 0x8103.
VlanIdList	If the port is a termination port, represents the filter lists VLAN in which the destination MAC address resides.

Configuring ACLs for mirroring

Configure the access control list (ACL) to mirror packets for an access control entry (ACE) that matches a particular packet.

Before you begin

- The ACL exists.

About this task

To modify an ACL parameter, double-click the parameter you wish to change. Change the value, and then click Apply. You cannot change a parameter that appears dimmed; in this case, delete the ACL, and then configure a new one.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Security > Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Double-click the parameter **MirrorVlanId** to configure mirroring to a destination VLAN.
5. Double-click the parameter **MirrorMltId** to configure mirroring to a destination MLT group.
6. Double-click the parameter **MirrorDstPortList** to configure mirroring to a destination port or ports.

ACL field descriptions

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


Name	Description
AcId	Specifies a unique identifier for the ACL from 1–2048.
Type	Specifies whether the ACL is VLAN- or port-based. Valid options are <ul style="list-style-type: none"> • inVlan • outVlan • inPort • outPort <p> Important: The inVlan and outVlan ACLs drop packets if you add a VLAN after ACE creation.</p>
Name	Specifies a descriptive user-defined name for the ACL.
VlanList	For inVlan and outVlan ACL types, specifies all VLANs to associate with the ACL.
PortList	For inPort and outPort ACL types, specifies the ports to associate with the ACL.
DefaultAction	Specifies the action taken when no ACEs in the ACL match. Valid options are deny and permit. Deny means the system drops the packets; permit means the system forwards packets. The default is permit.
ControlPktAction	Specifies the action for control packets, if you configure DefaultAction to deny. If DefaultAction is permit, this value is ignored.

Table continues...

Name	Description
State	Enables or disables all of the ACEs in the ACL. The default value is enable.
PktType	Indicates the packet type that this ACLI is applicable to. The default is IPv4.
IpfixState	Enables or disables the Internet Protocol Flow Information eXport (IPFIX) option on the ACL. Use IPFIX to monitor IP flows. The default is disabled.
MirrorVlanId	Configures mirroring to a destination VLAN.
MirrorMltId	Configures mirroring to a destination MLT group.
MirrorDstPortList	Configures mirroring to a destination port or ports.

Configuring ACEs for mirroring

Before you begin

- The ACL exists.
- The ACE exists.

About this task

Configure actions to use filters for flow mirroring. Use an ACE to define the mirroring actions the filter performs.

If you use the mirror action, ensure that you specify the mirroring destination: IP address, MLTs, ports, or VLANs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Security > Data Path**.
2. Click **Advanced Filters (ACE/ACLs)**.
3. Click the **ACL** tab.
4. Select the ACL for which to modify an ACE.
5. Click **ACE**.
6. Select an ACE, and then click **Action**.
7. Configure one of: **DstPortList**, **DstVlanId**, **DstMltId**, or **DstIp**.
8. Click **Apply**.

Action field descriptions

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

* Note:

The table lists the options for both Security ACEs and QoS ACEs. Dependent upon the ACE different options appear on the EDM interface.

Name	Description
AcId	Specifies the ACL ID from 1–2048
AcId	Specifies the ACE ID. Use ACE IDs 1–1000 for security rules. Use ACE IDs 1001–2000 for QoS rules.
Mode	Configures the action mode for security ACEs. The default value is deny.
MltIndex	<p>If you use this action, the ACE overrides the mlt-index chosen by the MLT algorithm for packets sent on MLT ports.</p> <p>The MLT index ranges from 0–16. If three ports exist in an MLT (for example, A, B, and C) and you specify an index of 6, the Virtual Services Platform 9000 applies the MOD function and chooses port C. If port C becomes nonoperational, the filtered packets exit the platform from port B.</p> <p>Multicast traffic does not support the MLT index. This variable is a security action. The ACE ID must be in the range of 1–1000.</p>
RemarkDscp	<p>Specifies the new Per-Hop Behavior (PHB) for matching packets: phbcs0, phbcs1, phbaf11, phbaf12, phbaf13, phbcs2, phbaf21, phbaf22, phbaf23, phbcs3, phbaf31, phbaf32, phbaf33, phbcs4, phbaf41, phbaf42, phbaf43, phbcs5, phbef, phbcs6, phbcs7.</p> <p>This action is a QoS action. The ACE ID must be in the range of 1001–2000.</p>
RemarkDot1Priority	<p>Specifies the new 802.1 priority bit for matching packets: zero, one, two, three, four, five, six, or seven.</p> <p>This action is a QoS action. The ACE ID must be in the range of 1001–2000.</p>
Police	Polices the packet according to the specified policy ID (0–16000). A policy must exist. This action is a QoS action. The ACE ID must be in the range of 1001–2000.
InternalQoS	This variable is a QoS action. The ACE ID must be in the range of 1001–2000. The default value is 1.
RedirectNextHop	<p>Specifies the next-hop IP address for redirect mode (a.b.c.d). This action is a security action. The ACE ID must be in the range of 1–1000.</p> <p>The default is 0.0.0.0.</p>

Table continues...

Name	Description
RedirectUnreach	Denies or permits packet dropping when the next-hop for the packet is unreachable. The default value is deny. This action is a security action. The ACE ID must be in the range of 1–1000.
IpfixState	Configures IPFIX metering using filters to use IPFIX on selected flows. An ACL can have multiple ACEs and each ACE can have an action enable. The default is disable.
Count	Enables the ability to count matching packets. Use this parameter with either a security or QoS ACE. The default is disabled.
Log	This action logs to the master CP module. Use this parameter with either a security or QoS ACE. The default is disabled.
CopytoPcap	This variable is a security action that sends a copy of the packet to the secondary CP module. The ACE ID must be in the range of 1–1000. The default is disabled.
DstPortList	Configures mirroring to a destination port or ports. This action is a security action. The ACE ID must be in the range of 1–1000.
DstVlanId	Configures mirroring to a destination VLAN. This action is a security action. The ACE ID must be in the range of 1–1000.
DstMltdId	Configures mirroring to a destination MLT group. This action is a security action. The ACE ID must be in the range of 1–1000.
DstIp	Configures Layer 3 mirroring. The destination must be an IP address {A.B.C.D}. For Layer 3 mirroring, the last hop in the path from the mirrored port to the remote mirroring destination should be on a VLAN by itself. Avaya recommends that you configure the remote mirrored port in its own VLAN at the last hop to prevent flooding. The hops between the mirror source port and the last hop can be on the same VLAN or on different VLANs and the hops between the mirror source port and the last hop can connect through bridging or routing.
DstIpDscp	Optionally, if you configure a destination IP address for mirroring, you can also configure the DSCP value. The DstIpDscp range is <0–63>. The default is 256 (disabled).

Table continues...

Name	Description
DstIpTtl	Optionally, if you configure a destination IP address for mirroring, you can also configure the time-to-live value. The DstIpTtl range is <2–255>. The default is 64.

Configuring PCAP globally

Use the Packet Capture Tool (PCAP) to capture packets for troubleshooting and security purposes. Configure PCAP globally to define how PCAP operates on the Avaya Virtual Services Platform 9000.

* Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Before you begin

- The secondary CPU is active.
- If you save to external storage, a Compact Flash (CF) card is installed.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **PCAP**.
3. Click the **PcapGlobal** tab.
4. Configure PCAP as required.
5. Click **Apply**.

PcapGlobal field descriptions

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

Name	Description
Enable	Enables or disables PCAP globally on the PCAP engine (secondary CPU). The default is disabled.
BufferWrap	Enables buffer wrap-around after the buffer is full. When enabled, PCAP continues to capture packets, otherwise, packet capturing stops.
ExtflashWrap	Select the ExtflashWrap checkbox to enable automatic overwriting of the file on the external flash or network during autosave. To prevent overwriting of the file on the external flash or network during autosave, deselect the ExtflashWrap checkbox.

Table continues...

Name	Description
FrameSize	Specifies the number of bytes of each packet that are captured. The default value is 64 bytes.
BufferSize	Specifies the amount of memory allocated for data. The default is 32 MB.
AutoSave	Saves data automatically after the buffer is full.
AutoSaveFileName	Specifies the name of the file in which packets are stored.
AutoSaveDevice	Specifies the device used to store the captured packets. If the device is network, you must enter an IP address.
AutoSaveNetworkIpAddress	Specifies the IP address of the remote host where the data must be stored. This field is valid only if the device is network.
CopyFileName	Specifies the file name to use when copying the PCAP file from the PCAP engine or an external storage device to a remote client (user local machine).

Configuring PCAP on a port

Configure PCAP on a port so that the port supports PCAP.

Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Before you begin

- If required, IP filters exist.
- If required, ACLs with a global action of mirror exist.

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 **General**.
4. Click the **PCAP** tab.
5. Select **Enable**.
6. Choose the PCAP mode.
7. Click **Apply**.

PCAP field descriptions

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

Name	Description
Enable	Enables or disables PCAP on the port. The default is disabled.
Mode	Configures the PCAP mode (tx, rx, or both). The default is rx mode.

Configuring PCAP filters

Use filters to narrow the scope of the types of packets to capture. Use these filters to match MAC and IP addresses, Differentiated Services Code Point (DSCP) and p-bit markings, VLAN IDs, and protocol types.

* Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **PCAP**.
3. Click the **PcapFilter** tab.
4. Click **Insert**.
5. Configure the filter as required.
6. Click **Insert**.

PcapFilter field descriptions

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

Name	Description
Id	Indicates the unique ID that represents the filter.
Enable	Enables or disables the filter. The default is disabled.
Action	<p>Determines the action taken by the filter:</p> <ul style="list-style-type: none"> • capture indicates that the packet is captured. This option is the default. • drop indicates that the packet is dropped. • trigger-off indicates that PCAP captures packets until one matches the criteria, and then disables the filter entry, and globally disables PCAP. • trigger-on indicates that PCAP captures a packet after it matches the criteria, and then disables the filter entry.

Table continues...


Name	Description
	<p> Important:</p> <p>Because the PCAP engine runs on the secondary CP module, the master CP module does not reflect the change in PCAP and PCAP capture-filter status if you use the action trigger-on or trigger-off. View capture filters on the secondary CP module to view the correct status. After PCAP disables the filter entry on the PCAP engine, if you view capture filters on the master CP module, the status appears as true (enabled), when it really is false (disabled). To activate PCAP and the PCAP capture filter again, you must reenable them on the master CP module.</p> <p>The default is capture.</p>
SrcMac	Specifies the source MAC address to match. The default is 6.
SrcMacMask	Specifies the source MAC address mask to specify an address range.
IsInverseSrcMac	Specifies the source MAC address inverse. If you select this variable, all MAC addresses other than the one specified are matched. The default is disabled.
DstMac	Specifies the destination MAC address.
DstMacMask	Specifies the destination MAC address mask to specify an address range. The default is 6.
IsInverseDstMac	Specifies the destination MAC address inverse. If you select this variable, all MAC addresses other than the one specified are matched. The default is disabled.
VlanId	Specifies the VLAN ID of the packet to match. The default is 0.
ToVlanId	Specifies the destination VLAN ID. Use to specify a range. The default is 0.
IsInverseVlanId	Specifies the VLAN ID inverse. If you select this variable, all VLAN IDs other than the one specified are matched. The default is disabled.
Pbit	Specifies the 802.1p-bit of the packet to match. The default is 0.
ToPbit	Specifies an 802.1p-bit range. The default is 0.
IsInversePbit	Specifies the p-bit inverse. If you select this variable, all p-bits other than the one specified are matched. The default is disabled.
PbitMatchZero	Instructs PCAP to consider 0 a valid p-bit value. Packets with a p-bit of 0 can be captured. Otherwise, 0 is considered a disable value. The default is disabled.
EtherType	Specifies the Ethertype of the packet to match. The default is 0.
ToEtherType	Specifies an Ethertype range. The default is 0.

Table continues...

Name	Description
IsInverseEtherType	Specifies the Ethertype inverse. If you select this variable, all Ethertypes other than the one specified are matched. The default is disabled.
SrcIpl	Specifies the source IP address of the packet to match.
ToSrcIpl	Specifies a source IP address range.
IsInverseSrcIpl	Specifies the source IP address inverse. If you select this variable, source IP addresses other than the one specified are matched. The default is disabled.
DstIpl	Specifies the destination IP address of the packet to match.
ToDstIpl	Specifies the destination IP address range.
IsInverseDstIpl	Specifies the destination IP address inverse. If you select this variable, all addresses other than the one specified are matched. The default is disabled.
Dscp	Specifies the DSCP of the packet to match.
ToDscp	Specifies a DSCP range.
IsInverseDscp	Specifies the DSCP inverse. If you select this variable, all DSCPs other than the one specified are matched. The default is disabled.
DscpMatchZero	Instructs PCAP to consider 0 a valid DSCP value. Packets with a DSCP of 0 can be captured. Otherwise, 0 is considered a disable value. The default is disabled.
ProtocolType	Specifies the protocol of the packet to match. The default is 0.
ToProtocolType	Specifies a protocol type range. The default is 0.
IsInverseProtocolType	Specifies the protocol type inverse. If you select this variable, all protocols other than the one specified are matched. The default is disabled.

Configuring advanced PCAP filters

Use advanced filters to match User Datagram Protocol (UDP) and TCP parameters, as well as to specify user-defined parameters.

* Note:

If a PCAP dump is in progress on the standby, you will encounter the following error if you try to make a PCAP configuration change:

```
Error: PCAP dump in progress. Command not allowed.
```

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **PCAP**.
3. Click the **PcapAdvancedFilter** tab.

4. Configure the filter as required.
5. Click **Apply**.

PcapAdvancedFilter field descriptions

Use the data in the following table to configure the **PcapAdvancedFilter** tab.

Name	Description
Id	Specifies the unique ID that represents the filter.
UdpPort	Specifies the UDP port of the packet to match. UdpPort can be one or a range of UDP port values. The default is 0.
ToUdpPort	Specifies a range of UDP ports. The default is 0.
IsInverseUdpPort	Indicates that all other values other than the specified range of UDP ports are matched. The default is disabled.
TcpPort	Specifies the TCP port of the packet to match. The default is 0.
ToTcpPort	Specifies a range of TCP ports. The default is 0.
IsInverseTcpPort	Indicates that all other values other than the specified range of TCP ports are matched. The default is disabled.
UserDefinedData	Specifies the user-defined data to match. The default is 0.
UserDefinedDataSize	Specifies the length of user-defined data. The default is 0.
UserDefinedOffset	Specifies the offset from which the match must start. The default is 0.
IsInverseUserDefined	Indicates that all data other than the specified user-defined data is matched. The default is disabled.
Timer	Specifies that PCAP is invoked after the first packet matches and stops after a configured value of time. After the timer starts, the filter is disabled. After the PCAP engine receives a matching packet, it captures all packets for the duration of the timer, and then disables PCAP globally. Specify the timer value in milliseconds (ms). This option is active only if the filter action is trigger-on. The default value is 0.
PacketCount	Stops PCAP capturing after capturing the specified value of packets. This action is similar to the refresh-timer option; once it is invoked, the filter is disabled. This variable is active only if the filter action is trigger-on. To delete this option, configure it to 0. The default value is 0.
RefreshTimer	Starts or restarts the timer. After the PCAP engine receives a matching packet, it disables the capture filter. If the PCAP engine does not receive another matching packet within the specified time, PCAP is disabled globally. The timer restarts every time the PCAP engine receives a packet, until the timer expires. Specify the value in ms. This variable is active only if the filter action is trigger-on. To delete this option, configure it to 0. The default value is 0.

Running a ping test

About this task

Use ping to determine if an entity is reachable.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Ping Control** tab.
4. Click **Insert**.
5. In the **OwnerIndex** box, type the owner index.
6. In the **TestName** box, type the name of the test.
7. In the **TargetAddress** box, type the host IP address.
8. From the **AdminStatus** options, select **enabled**.
9. In the remainder of the option boxes, type the desired values.
10. Click **Insert**.
11. Select an entry, and then click **Start**.
Let the test run for several seconds.
12. Select an entry, and then click **Stop**.
13. View the Ping results.

Ping Control field descriptions

Use the data in the following table to use the **Ping Control** tab.

Name	Description
OwnerIndex	Provides access control by a security administrator using the View-Based Access Control Model (VACM) for tables in which multiple users need to independently create or modify entries. This is a string of up to 32 characters.
TestName	Specifies the name of the ping test.
TargetAddressType	Specifies the type of host address to use at a remote host to perform a ping operation.
TargetAddress	Specifies the host address to use at a remote host to perform a ping operation.
DataSize	Specifies the size of the data portion (in octets) to transmit in a ping operation. The default is 0.

Table continues...

Name	Description
TimeOut	Specifies the timeout value, in seconds, for a remote ping operation. The default is 3 seconds.
ProbeCount	Specifies the number of times to perform a ping operation at a remote host. The default is 1.
AdminStatus	Specifies the state of the ping control entry: enabled or disabled. The default is disabled.
DataFill	Determines the data portion of a probe packet.
Frequency	Specifies the number of seconds to wait before repeating a ping test. The default is 0.
MaxRows	Specifies the maximum number of entries allowed in the PingProbeHistory table. The default is 50.
StorageType	Specifies the storage type for this row.
TrapGeneration	<p>Specifies when to generate a notification. The options are:</p> <ul style="list-style-type: none"> • ProbeFailure—Generates a PingProbeFailed notification subject to the value of TrapProbeFailureFilter. The object TrapProbeFailureFilter can specify the number of successive probe failures that are required before a pingProbeFailed notification is generated. • TestFailure—Generates a PingTestFailed notification. The object TrapTestFailureFilter can determine the number of probe failures that signal when a test fails. • TestCompletion—Generates a PingTestCompleted notification. <p>The value of this object defaults to zero, indicating that none of the preceding options have been selected.</p>
TrapProbeFailureFilter	Specifies the number of successive probe failures that are required before a pingProbeFailed notification is generated. The default is 1.
TrapTestFailureFilter	Determines the number of probe failures that signal when a test fails. The default is 1.
Type	Selects or reports the implementation method used to calculate ping response time. The default is pingIcmpEcho.
Descr	Describes the remote ping test. The default is 0x00.
SourceAddressType	Specifies the type of source address used at a remote host when performing a ping operation. The default is IPv4.
SourceAddress	Specifies the IP address (a.b.c.d) as the source address in outgoing probe packets.
IfIndex	Setting this object to the ifIndex of an interface, prior to starting a remote ping operation, directs the ping probes to be transmitted over the specified interface. The default is 0.

Table continues...

Name	Description
ByPassRouteTable	Enables (optionally) the bypassing of the route table. The default is disabled.
DSField	Specifies the value to store in the Differentiated Services (DS) field in the IP packet used to encapsulate the ping probe. The default is 0.

Viewing ping results

About this task

View ping results to view performance-related data.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Ping Control** tab.
4. Select a ping test entry.
5. Click **Ping Result**.

Ping Result field descriptions

Use the data in the following table to use the **Ping Result** tab.

Name	Description
OwnerIndex	Specifies the ping test owner.
TestName	Specifies the test name.
OperStatus	Indicates the operational status of the test. The default is disabled.
IpTargetAddressType	Specifies the IP address type of the target. The default is unknown.
IpTargetAddress	Specifies the IP address of the target.
MinRtt	Specifies the minimum ping round-trip-time (RTT) received. A value of 0 means that no RTT is received.
MaxRtt	Specifies the maximum ping RTT received. A value of 0 means that no RTT is received.
AverageRtt	Specifies the current average ping RTT.
ProbeResponses	Specifies the number of responses to probes.
SentProbes	Specifies the number of sent probes.
RttSumOfSquares	Specifies the sum of squares of RTT for all probes received.
LastGoodProbe	Specifies the date and time when the last response is received for a probe.

Viewing ping probe history

About this task

View the ping probe history to view the history of ping tests performed by the switch.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **Ping/Trace Route**.
3. Select a ping entry.
4. Click **Ping Probe History**.

Ping Probe History field descriptions

Use the data in the following table to use the **Ping Probe History** tab.

Name	Description
OwnerIndex	Specifies the owner index.
TestName	Indicates the name given to the test.
Index	Specifies the index number.
Response	Indicates the amount of time, measured in milliseconds, between request (probe) and response, or when the request timed out. Response is reported as 0 when it is not possible to transmit a probe.
Status	Indicates the status of the response; the result of a particular probe done by a remote host.
LastRC	Indicates the last implementation-method-specific reply code (RC) received. If ICMP echo is used, then a successful probe ends when an ICMP response is received that contains the code ICMP_ECHOREPLY(0).
Time	Indicates the timestamp for this probe result.

Running a traceroute test

Run a traceroute test to determine the route packets take through a network to a destination.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Trace Route Control** tab.
4. Click **Insert**.

5. Configure the instance as required.
6. Click **Insert**.
7. Select an entry, and then click **Start**.
Let the test run for several seconds.
8. Select an entry, and then click **Stop**.
9. View the traceroute test results.

Trace Route Control field descriptions

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

Name	Description
OwnerIndex	Provides access control by a security administrator using the VACM for tables in which multiple users need to independently create or modify entries.
TestName	Specifies the name of the traceroute test.
TargetAddressType	Specifies the type of host address, either IPv4 or IPv6, to use on the traceroute request at the remote host. The default is IPv4.
TargetAddress	Specifies the host address used on the traceroute request at the remote host.
ByPassRouteTable	Enables bypassing of the route table. If you enable this variable, the remote host bypasses the normal routing tables and sends directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. You can use this variable to perform the traceroute operation to a local host through an interface that has no route defined. The default is disabled.
DataSize	Specifies the size of the data portion of a traceroute request in octets. The default is 1.
TimeOut	Specifies the timeout value, in seconds, for a traceroute request. The default is 3.
ProbesPerHop	Specifies the number of times to reissue a traceroute request with the same time-to-live (TTL) value. The default is 3.
Port	Specifies the UDP port to which to send the traceroute request. Specify a port that is not in use at the destination (target) host. The default is the IANA assigned port 33434.
MaxTtl	Specifies the maximum time-to-live from 1–255. The default is 30.
DSField	Specifies the value to store in the Differentiated Services (DS) field in the IP packet used to encapsulate the traceroute probe. The default is 0.
SourceAddressType	Specifies the type of the source address to use for a remote host.
SourceAddress	Uses the specified IP address (which must be an IP number, not a hostname) as the source address in outgoing probe packets.

Table continues...

Name	Description
IfIndex	Directs the traceroute probes to be transmitted over the specified interface. The default is 0.
MiscOptions	Enables an application to specify implementation-dependent options.
MaxFailures	Indicates the maximum number of consecutive timeouts allowed before terminating a remote traceroute request. The default is 5.
DontFragment	Enables setting of the do not fragment (DF) flag in the IP header for a probe. The default is disabled.
InitialTtl	Specifies the initial time-to-live (TTL) value to use. The default is 1.
Frequency	Specifies the number of seconds to wait before repeating a traceroute test as defined by the value of the various objects in the corresponding row. The default is 0.
StorageType	Specifies the storage type for this row.
AdminStatus	Specifies the desired state for TraceRouteCtlEntry. The options are enabled or disabled. The default is disabled.
MaxRows	Specifies the maximum number of entries allowed in the TraceRouteProbeHistoryTable. The default is 50.
TrapGeneration	Determines when to generate a notification for this entry. The options are: <ul style="list-style-type: none"> • pathChange —Generates a TraceRoutePathChange notification after the current path varies from a previously determined path. • testFailure —Generates a TraceRouteTestFailed notification after the full path to a target cannot be determined. • testCompletion —Generates a TraceRouteTestCompleted notification after the path to a target has been determined.
Descr	Describes the remote traceroute test.
CreateHopsEntries	Stores the current path for a traceroute test in the TraceRouteHopsTable on an individual hop basis when the value of this object is true. The default is false.
Type	Reports or selects the implementation method to use for performing a traceroute operation.

Viewing traceroute results

About this task

View traceroute results to view performance-related data.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **Ping/Trace Route**.

3. Click the **Trace Route Control** tab.
4. Select a traceroute entry.
5. Click **Trace Route Result**.

Trace Route Result field descriptions

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

Name	Description
OwnerIndex	Specifies the index of the owner.
TestName	Specifies the name of the test.
OperStatus	Specifies the operational status of the test. The default is disabled.
CurHopCount	Specifies the current count of hops.
CurProbeCount	Specifies the current count of probes.
IpTgtAddressType	Specifies the IP target address type
IpTgtAddr	Specifies the IP target address.
TestAttempts	Specifies the number of test attempts.
TestSuccesses	Specifies the number of successful test attempts.
LastGoodPath	Specifies the date and time when the last response is received for a probe.

Viewing the traceroute history

About this task

View the traceroute history to view the history of traceroute tests performed by the switch.

The traceroute probe history contains probe information for the hops in the routing path.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **Ping/Trace Route**.
3. Click the **Trace Route Control** tab.
4. Select an entry.
5. Click **Trace Route Probe History**.

Route Probe History field descriptions

Use the data in the following table to use the **Trace Route Probe History** tab.

Name	Description
OwnerIndex	Identifies the Trace Route entry to which a probe result belongs.

Table continues...

Name	Description
TestName	Specifies the test name.
Index	Specifies the Index.
HopIndex	Indicates for which hop in a traceroute path the probe results are intended.
ProbeIndex	Specifies the index of a probe for a particular hop in a traceroute path.
HAddrType	Specifies the IP address type of the hop to which this probe belongs.
HAddr	Specifies the IP address of the hop to which this probe belongs.
Response	Specifies the cumulative results at any time.
Status	Specifies the status of the probe.
LastRC	When a new entry is added, the old entry is purged if the total number of entries exceeds the specified maximum number of entries in the Control Table Entry.
Time	Specifies the response time of the probe.

Chapter 8: Layer 1 troubleshooting

Use this section to help you troubleshoot Layer 1 (physical layer) problems.

Troubleshooting fiber optic links

About this task

You can troubleshoot fiber optic links to ensure that the optical transmitters and receivers operate correctly, and to determine if a receiver is saturated, or does not receive enough power.

To troubleshoot optical links and devices, you can use Digital Diagnostic Monitoring (DDM), as well as published optical specifications.

For more information about quad small form factor pluggable plus (QSFP+), SFP+s, and SFP transceivers and SFP+ transceivers, see *Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000*, NN46250-305.

Procedure

1. Measure the QSFP+, SFP+s, and SFP transmit power.
2. Compare the measured transmit power with the specified launch power.

The values are similar. If the measured power is far below the specified value, a faulty transmitter is a possible cause.

3. Compare the measured transmit power for the near-end optical device to the measured transmit power for the far-end device.

Large differences can mean that the optical devices are mismatched (that is, -SX versus -LX).

4. Measure the receive power at each end of the link.
5. Compare the receive power to the transmit power.

- For short fiber links, the transmit and received power are similar (after taking into consideration connection losses).
- For long fiber links, the transmit and received power are similar (after taking into consideration connection losses and fiber attenuation).

Large differences can mean a damaged fiber or dirty or faulty connectors. Large differences can also mean that the link does not use the right type of fiber (single mode or multimode). If

the receiver power is measured to be zero, and the link worked previously, it is probable that the far-end transmitter is not operating or the fiber is broken.

6. Compare the measured receive power for the near-end optical device to the measured receive power for the far-end device.

Large differences can mean that the optical devices are mismatched (that is, -SX versus -LX). If optical devices are mismatched, the receiver can be saturated (overdriven).

7. If a receiver is saturated but still operable, install a suitable attenuator.

For long-haul optical devices, the receive power must be significantly less than the transmit power.

8. To help debug the link, loop back the local transmit and receive ports, and use the DDM parameters to help determine the fault.

Resetting a QSFP+ transceiver

Reset a transceiver to simulate removal and reinsertion of the transceiver, which can be helpful in troubleshooting. For example, if authentication of the transceiver fails but you believe the transceiver is a qualified Avaya part, you can reset the transceiver to begin the authentication process again.

About this task

Resetting the transceiver stops traffic and triggers log messages similar to the removal and insertion of the transceiver.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Reset the transceiver:

```
pluggable-optical-module reset {slot/port}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#pluggable-optical-module reset 3/12

CP1 [11/27/14 14:18:43.503] 0x00004726 00000000 GlobalRouter SNMP INFO SPBM detected adj
UP on Port10/1, neighbor 00bb.0000.8100 (VSP8000-1)
CP1 [11/27/14 14:18:31.511] 0x000ec5cf 04f00001.640 DYNAMIC CLEAR GlobalRouter LACP INFO
trapSendVlaccpLinkUp: Vlaccp Link 10/1 is up
CP2 [11/27/14 14:18:31.078] 0x00000026 00000000 GlobalRouter SW INFO Port 10/1 is a
trunk port
CP1 [11/27/14 14:18:31.068] 0x00000026 00000000 GlobalRouter SW INFO Port 10/1 is a
trunk port
CP1 [11/27/14 14:18:31.068] 0x0000c5ec 00300001.640 DYNAMIC CLEAR GlobalRouter HW INFO
Link Up(10/1)
```

```

CP2 [11/27/14 14:18:30.836] 0x00000026 00000000 GlobalRouter SW INFO Port 10/1 is a
trunk port
CP1 [11/27/14 14:18:30.823] 0x00004726 00000000 GlobalRouter SNMP INFO SPBM detected adj
DOWN on Port10/1, neighbor 00bb.0000.8100 (VSP8000-1)
CP1 [11/27/14 14:18:30.822] 0x000ec5ce 04f00001.640 DYNAMIC SET GlobalRouter LACP
WARNING trapSendVlacpLinkDown: Vlacp Link 10/1 is down
CP1 [11/27/14 14:18:30.818] 0x00000026 00000000 GlobalRouter SW INFO Port 10/1 is a
trunk port
CP1 [11/27/14 14:18:30.818] 0x0000c5e7 00300001.640 DYNAMIC SET GlobalRouter HW INFO
Link Down(10/1)
CP1 [11/27/14 14:18:30.815] 0x002c0600 00000000 GlobalRouter CLILOG INFO 1655 TELNET:
135.64.95.188 rwa pluggable-optical-module reset 10/1
CP1 [11/27/14 14:18:16.691] 0x002c0600 00000000 GlobalRouter CLILOG INFO 1654 TELNET:
135.64.95.188 rwa show pluggable-optical-modules basic

```

Variable definitions

Use the data in the following table to use the **pluggable-optical-module reset** command.

Variable	Value
slot/port	Specifies the slot and port location of the QSFP+ transceiver to reset.

Chapter 9: Layer 2 troubleshooting

Use this section to help you troubleshoot virtual LAN (VLAN), link aggregation, and MultiLink Trunking (MLT) problems.

Troubleshooting IST failure

About this task

When you use interswitch trunk (IST) links, all critical network traffic runs on this link. If the IST fails, network protocols, for example, Routing Information Protocol (RIP), Virtual Router Redundancy Protocol (VRRP), Open Shortest Path First (OSPF), and Virtual Link Aggregation Control Protocol (VLACP), go up and down and eventually cause a network outage.

Possible reasons for IST failure are:

- The IST is disabled.
- The IST uses an incorrect peer IP address.
- The MLT does not use the proper ports.
- MultiLink Trunking (MLT) ports are down.
- The IST VLAN ID is different from that of the peer.
- The VLACP port state is down on all IST ports, if VLACP enabled.
- The ARP entry for the IST peer IP address is missing.
- One side runs Release 3.2 software and the other side runs a later release, for example, Release 3.3.

Procedure

1. Check the IST configuration:

```
show ist mlt
```

2. If the configuration is incorrect, make the correction.

3. Check the MLT configuration:

```
show mlt <1-512>
```

4. Verify that the correct ports are members of the MLTs, and that the trunk is enabled.

5. Check interface status:


```
show ip interface
```

6. Verify that, for each interface, the port state is enabled, and the operational state is up. If the states are not enabled and up, try to disable, and then reenables the port. If the ports do not come up, the cause can be a physical layer issue.

7. Check the VLACP port state for the IST ports, if VLACP on those ports is enabled.

```
show vlacp interface gigabitethernet {slot/port[-slot/port][, ...]}
```

If the VLACP port state is down, disable and reenables VLACP on those ports.

8. Check the ARP entry for the IST peer IP address.

```
show ip arp {A.B.C.D}
```

If the ARP entry does not exist, shutdown and reenables the IST ports or disable and reenables the IST.

9. Check the software version.

```
show software [verbose]
```

If the two sides run different versions, upgrade the software.

Example

Check the IST configuration:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#show ist mlt
```

```
=====
Mlt IST Info
=====
```

MLT ID	PEER-IP ADDRESS	VLAN ID	ENABLE IST	IST STATUS
19	190.0.0.66	1900	true	up

```
-----
NEGOTIATED DIALECT          IST STATE          MASTER/SLAVE
-----
v4.0          Up          Slave
-----
```

```
VSP-9012:1(config)#show mlt
```

```
=====
Mlt Info
=====
```

MLTID	IFINDEX	NAME	PORT TYPE	MLT ADMIN	MLT CURRENT	PORT MEMBERS	VLAN IDS
2	6145	MLT-2	access	smlt	smlt	6/15,10/15	10
3	6146	MLT-3	access	smlt	norm		10
19	6162	MLT-19	trunk	ist	ist	6/17	10 1900

All 3 out of 3 Total Num of mlt displayed

MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER

Layer 2 troubleshooting

```

-----
2      6145    10/15      disable   down
3      6146    null       disable   down
19     6162    6/17      disable   down

All 3 out of 3 Total Num of mlt displayed

                                           WHICH PORTS

--More-- (q = quit)

VSP-9012:1>show ip interface

=====
                          IP Interface - GlobalRouter
=====
INTERFACE      IP          NET          BCASTADDR  REASM      VLAN  BROUTER
ADDRESS        MASK        FORMAT      MAXSIZE    ID         PORT
-----
Vlan3998      30.10.10.1  255.0.0.0   ones       1500      3998  false
Vlan4000      10.10.10.1  255.0.0.0   ones       1500      4000  false

All 2 out of 2 Total Num of IP interfaces displayed

VSP-9012:1#show vlacp interface gigabitethernet

=====
                          VLACP Information
=====
INDEX  ADMIN   OPER    PORT   FAST   SLOW   TIMEOUT  TIMEOUT  ETHER   MAC
      ENABLED  ENABLED STATE  TIME   TIME   TIME     SCALE   TYPE   ADDR
-----
4/1   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/2   true   true    UP     200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/3   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/4   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/5   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/6   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/7   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/8   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/9   false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/10  false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/11  false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/12  false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/13  false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/14  true   true    UP     200    30000  long     3       0x8103  01:80:c2:
00:11:00
4/15  false  false   DOWN   200    30000  long     3       0x8103  01:80:c2:
00:11:00

```

```

4/16 false false DOWN 200 30000 long 3 0x8103 01:80:c2:
00:11:00

--More-- (q = quit)

VSP-9012:1#show ip arp 31.30.30.32

=====
                        IP Arp - GlobalRouter
=====
IP_ADDRESS      MAC_ADDRESS      VLAN      PORT TYPE      TTL(10 Sec)
-----
31.30.30.32     00:14:c7:5f:42:00 3        MLT 1    DYNAMIC 1177
=====

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

1 out of 51 ARP entries displayed

VSP-9012:1(config)#show software verbose

=====
                        software releases in /intflash/release/
=====
Release          Added Time          Activated Time          Committed Time          Committed Type
-----
3.1.0.0.GA      -----
3.2.0.0.GA      -----
3.3.0.0.GA      -----
(Backup Release)
3.4.0.0.GA      2013-09-27 13:55:56 2013-09-27 14:00:00 2013-09-27 14:13:00 Auto
(Primary Release)

-----

Auto Commit      : enabled
Commit Timeout   : 10 minutes

```

Variable definitions

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

Table 29: Variable definitions

Variable	Value
1-512	Specifies the MLT ID.
error [collision][main]	Displays MLT statistic error information. The collision parameter displays collision error information. The main parameter displays MLT error main information.
stats	Displays MLT statistics.

Use the data in the following table to use the `show vlacp interface gigabitethernet` command.

Table 30: Variable definitions

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

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

Table 31: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the IP address for the IST peer.

Troubleshooting BPDU Filtering

The following procedures provide information to troubleshoot issues with Bridge Protocol Data Unit (BPDU) Filtering.

No packets received on the port

For BPDU Filtering to work on a port, the port must receive BPDU packets. Perform the following procedure to troubleshoot cases when the port does not receive packets.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Show the BPDU Filtering status for the port:

```
show spanning-tree bpdu-filtering {slot/port[-slot/port][,...]}
```

3. Use the following command to verify that the port receives packets:

```
show interface gigabitEthernet statistics verbose {slot/port[-slot/port][,...]}
```

4. Verify that the remote port is sending packets:

```
show spanning-tree {mstp|rstp} port role [{slot/port[-slot/port][,...]}]
```

```
show spanning-tree {mstp|rstp} port statistics [{slot/port[-slot/
port][, ...]]}
```

Example

The following example shows that BPDU Filtering is enabled for port 4/1. The BPDU Filter administrative state for the port is enabled but the timer counter is 0.

Port 4/1 receives packets. The remote port is disabled and does not send BPDU packets.

```
VSP-9012:1>enable
VSP-9012:1#show spanning-tree bpdu-filtering 4/1
Port MLTID Admin Oper Link LinkTrap Timeout TimerCount BpduFiltering
-----
4/1 Enable Up Up Enabled 200 0 Enabled
VSP-9012:1#show interface gigabitEthernet statistics verbose 4/1
=====
Port Stats Interface Extended
=====
PORT_NUM IN_UNICST OUT_UNICST IN_MULTICST OUT_MULTICST IN_BRDCST OUT_BRDCST
IN_LSM OUT_LSM
-----
4/1 201 0 160062 60943 4 72
0 0
VSP-9012:1#show spanning-tree mstp port role 10/11
=====
CIST Port Roles and States
=====
Port-Index Port-Role Port-State PortSTPStatus PortOperStatus
-----
12/11 Disabled Forwarding Disabled Disabled
VSP-9012:1#show spanning-tree mstp|rstp port statistics 10/11
=====
MSTP Cist Port Statistics
=====
Port Number : 12/11
Cist Port Fwd Transitions : 0
Cist Port Rx MST BPDUs Count : 0
Cist Port Rx RST BPDUs Count : 0
Cist Port Rx Config BPDUs Count : 0
Cist Port Rx TCN BPDUs Count : 0
Cist Port Tx MST BPDUs Count : 0
Cist Port Tx RST BPDUs Count : 0
Cist Port Tx Config BPDUs Count : 0
Cist Port Tx TCN BPDUs Count : 0
Cist Port Invalid MSTP BPDUs Rx : 0
Cist Port Invalid RST BPDUs Rx : 0
Cist Port Invalid Config BPDUs Rx : 0
Cist Port Invalid TCN BPDUs Rx : 0
Cist Port Proto Migr Count : 0
```

Variable definitions

Use the data in the following table to use the **show spanning-tree bpdu-filtering** command.

Variable	Value
{slot/port[-slot/port][, ...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Use the data in the following table to use the **show interface gigabitEthernet statistics verbose** command.

Variable	Value
{slot/port[-slot/port][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Use the data in the following table to use the **show spanning-tree** command.

Variable	Value
{mstp rstp}	Specifies the spanning tree protocol.
{slot/port[-slot/port][,....]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

SNMP trap not received

Perform the following procedure to troubleshoot issues in which an SNMP trap is not received.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Show the BPDU Filtering status for the port:

```
show spanning-tree bpdu-filtering {slot/port[-slot/port][,....]}
```

3. Configure the correct SNMP target information:

```
snmp-server host WORD<1-256> [port <1-65535>] v3 {noAuthNoPriv|
authNoPriv|authPriv WORD<1-32> [inform [timeout <1-2147483647>]
[retries <0-255>]] [filter WORD<1-32>]}
```

Example

In the following example, BPDU filtering is enabled on port 4/1, BPDU packets are received, port 4/1 is disabled, and the TimerCount is incrementing, but no SNMP trap is ever received.

```
VSP-9012:1>enable
VSP-9012:1#show spanning-tree bpdu-filtering 4/1
Port MLTID Admin Oper Link LinkTrap Timeout TimerCount BpduFiltering
-----
4/1 Disable Down Down Enabled 200 116 Enabled
```

Variable definitions

Use the data in the following table to use the **show spanning-tree** command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Use the data in the following table to use the `snmp-server host` command.

Variable	Value
filter <i>WORD</i> <1-32>	Specifies a filter profile name.
host <i>WORD</i> <1-256>	Specifies the IPv4 or IPv6 host address
inform [timeout <1-2147483647>]	Specifies the notify type. The optional timeout parameter configures the timeout value, which specifies the time to wait for a reply before resending the inform message. Time is specified in centiseconds
noAuthNoPriv authNoPriv authPriv <i>WORD</i> <1-32>	Specifies the security level.
port <1-65535>	Specifies the port number that will be set as the destination port at the UDP level in the trap packet.
retries <0-255>	Specifies the number of packets to be sent if no reply is received.
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

Chapter 10: Connectivity Fault Management

Use Connectivity Fault Management (CFM) to troubleshoot Shortest Path Bridging MAC (SPBM) cloud and the customer domain.

CFM fundamentals

In a network, you need the ability to isolate a connectivity fault to correct it. When the network consists of an SPBM cloud as well as a customer domain, Connectivity Fault Management (CFM) helps determine where the problem exists to debug connectivity issues and isolate faults. By allowing CFM to break a network into sections using MEPs and MIPs, you can determine where the problem lies.

Typically the backbone nodes only learn Backbone MAC (B-MAC) addresses, while only the appropriate Backbone Edge Bridges (BEBs), which terminate the virtual services networks (VSN), learn the Customer MAC (C-MAC) addresses. As such, the nodes within the SPBM backbone have no knowledge of the C-MAC or IP addresses used within the Virtual Services Networks (VSNs) and only need to provide reachability to the B-MAC addresses within the backbone.

CFM divides or separates a network into administrative domains called Maintenance Domains (MD). CFM further subdivides each MD into logical groupings called Maintenance Associations (MA). A single MD can contain several MAs.

Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Two types of MP exist:

- Maintenance End Point (MEP)
- Maintenance Intermediate Point (MIP)

By allowing CFM to break a network into sections using MEPs and MIPs, the user can determine where in the network the problem exists.

You can explicitly configure MDs, MAs, MEPs and MIPs and associate them with multiple VLANs or you can use autogenerated CFM commands that create a MEP and MIP at a specified level for every SPBM B-VLAN or C-MAC C-VLAN. If you choose to autogenerated CFM commands, the VSP 9000 creates the MD, MA and MEP ID used for each MEP.

*** Note:**

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN. VSP 9000 only supports one MEP or MIP on the SPBM B-VLAN.

CFM provides loopback messages (LBM), which act like ping. CFM also provides linktrace messages (LTM), which act like traceroute. As a result, you can debug Layer 2 with CFM. The Virtual Services Platform 9000 wraps the CFM LBM and LTM into easier commands, namely `l2 ping` and `l2 traceroute` respectively. The l2 commands only require a VLAN and destination target MAC address to use.

The MEPs and MIPs that you configure for SPBM B-VLANs do not respond to CFM messages sent from C-MAC VLANs because the VLAN used is not the same and packet encapsulation is different. You must use autogenerated CFM MEP and MIP level for every C-VLAN on the chassis. You can use either explicitly configured or autogenerated CFM MEP and MIP for SPBM B-VLANs.

CFM is based on the IEEE 802.1ag standard. To support troubleshooting of the SPBM cloud, Virtual Services Platform 9000 supports a subset of CFM functionality. CFM is the standard for Layer 2 ping, Layer 2 traceroute, and the end-to-end connectivity check of the Ethernet network.

On Virtual Services Platform 9000, CFM is implemented to support Loopback Messages (LBM), and Linktrace Messages (LTM). Messages are sent between Maintenance Points (MP) in the system. Continuity Check Messages (CCM) are not required or supported in the current release.

Autogenerated CFM and explicitly configured CFM

VSP 9000 simplifies CFM configuration with autogenerated CFM. With autogenerated CFM, you use the commands `cfm spbm enable` and `cfm cmac enable` and VSP 9000 creates default MD, MA, MEPs, and MIPs for SPBM B-VLANs and C-VLANs respectively.

- For SPBM B-VLANs, VSP 9000 provides two methods to configure CFM: autogenerated and explicitly configured. You cannot use both.
- For C-VLANs, you can only use autogenerated CFM.

Autogenerated CFM

You can use autogenerated CFM at a global level to create a MEP and a MIP at a specified level for every SPBM B-VLAN and C-VLAN on the chassis. If you use autogenerated CFM commands, you do not have to configure explicit MDs, MAs, MEPs, or MIPs, and associate them with multiple VLANs.

If you do not want to use autogenerated CFM commands, you can choose to configure explicit MDs, MAs, MEPs, and MIPs for SPBM B-VLANs. However, you cannot use both an autogenerated CFM configuration and an explicit CFM configuration together.

*** Note:**

Previous explicit CFM configurations of MDs, MAs, and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you

must first remove the existing MEP and MIP on the SPBM B-VLANs. VSP 9000 only supports one type of MEP or MIP for each SPBM B-VLAN.

For autogenerated CFM configuration information for ACLI see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 197.
- [Configuring autogenerated CFM on C-VLANs](#) on page 198.

For autogenerated CFM configuration information for EDM see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 224.
- [Configuring autogenerated CFM on C-VLANs](#) on page 226.

Explicitly configured CFM

If you choose to explicitly configure CFM, you must configure an MD, MA, MEPS, and MIPs.

For explicit configuration information for ACLI see:

- [Configuring CFM MD](#) on page 201.
- [Configuring CFM MA](#) on page 202.
- [Assigning a MEP MIP level to an SPBM B-VLAN](#) on page 203.
- [Configuring CFM MEP](#) on page 205.

For explicit configuration information for EDM see:

- [Configuring CFM MD](#) on page 228.
- [Configuring CFM MA](#) on page 228.
- [Configuring CFM MEP](#) on page 229.
- [Configuring CFM nodal MEP](#) on page 231.

Using CFM

For SPBM B-VLANs, the autogenerated MEPS and MIPs respond to `12 ping`, `12 traceroute`, and `12 tracertree` in the same manner as the MEPS and MIPs created explicitly. For C-VLANs, the autogenerated MEPS and MIPs respond to `12 ping` and `12 traceroute`, but not to `12 tracertree` because no multicast trees exist on C-VLANs. The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPS.

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to `12 ping` and `12 traceroute` requests.

Customer VLAN vs. SPBM B-VLAN configurations

CFM breaks the network into sections, called MEPS, so you can determine exactly where the problem exists.

The MEPS and MIPs configured for SPBM B-VLANs do not respond to CFM messages sent by C-VLANs.

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC)

addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a B-MAC destination address (BMAC-DA) and a B-MAC source address (BMAC-SA). In SPBM, each node populates its forwarding database (FDB) with the B-MAC information derived from the IS-IS shortest path tree calculations.

Typically the SPBM Backbone Core Bridges (BCBs) in the SPBM cloud only learn the B-MAC addresses. The Backbone Edge Bridges (BEBs) know the Customer MACs on the appropriate BEBs that terminate the virtual services networks (VSNs). As such, the nodes within the SPBM cloud have no knowledge of the C-MAC addresses in the VSNs.

! Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For C-VLANs, you have to trigger an `12 ping` to learn the C-MAC address.
- For B-VLANs, you do not have to trigger an `12 ping` to learn the C-MAC address because IS-IS populates the MAC addresses in the FDB table.

In both cases, linktrace traces the path up to the closest device to that MAC address that supports CFM in the SPBM cloud.

C-VLANs source addresses

CFM uses either the VLAN MAC or the CFM C-MAC for the BMAC-SA for the C-VLANs. The CFM C-MAC is the value of the management base MAC, which ends in 0x3c0. The system creates the VLAN MAC after a user adds an IP address to a VLAN.

If a VLAN has a MAC address, the system uses the VLAN MAC as the BMAC-SA by default. If a VLAN does not have a MAC address, the system uses the CFM C-MAC for the BMAC-SA. You may also configure the system to use the CFM C-MAC, even if a VLAN MAC exists.

Maintenance Domain (MD)

A Maintenance Domain (MD) is the part of a network that is controlled by a single administrator. For example, a customer can engage the services of a service provider, who, in turn, can engage the services of several operators. In this scenario, there can be one MD associated with the customer, one MD associated with the service provider, and one MD associated with each of the operators.

You assign one of the following eight levels to the MD:

- 0–2 (operator levels)
- 3–4 (provider levels)
- 5–7 (customer levels)

The levels separate MDs from each other and provide different areas of functionality to different devices using the network. An MD is characterized by a level and an MD name (optional).

A single MD can contain several Maintenance Associations (MA).

Maintenance Association (MA)

An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

The following figure shows MD level assignment in accordance with the 802.1ag standard. As shown in the figure, MIPs can be associated with MEPs. However, MIPs can also function independently of MEPs.




Figure 4: MD level assignment

Maintenance endpoints (MEP)

A Maintenance Endpoint (MEP) defines the end of a link and a Maintenance Intermediate Point is a point in the middle of the network.

A Maintenance Endpoint (MEP) represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA. MEP functionality can be divided into the following functions:

- Fault Detection
- Fault Verification
- Fault Isolation
- Fault Notification

Fault detection and notification are achieved through the use of Continuity Check Messages (CCM). CCM messages are not supported in the current release.

Maintenance domain intermediate points (MIP)

MIPs do not initialize any CFM messages. MIPs passively receive CFM messages, process the messages received and respond back to the originating MEP. By responding to received CFM messages, MIPs can support discovery of hop-by-hop path among MEPs, allow connection failures to be isolated to smaller segments of the network to help discover location of faults along the paths. MIPs can be created independent of MEPs. MIP functionality can be summarized as:

- Respond to Loopback (ping) messages at the same level as itself and addressed to it.
- Respond to Linktrace (traceroute) messages.
- Forward Linktrace messages after decrementing the TTL.

Fault verification

Fault verification is achieved through the use of Loopback Messages (LBM). An LBM is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a MEP or Maintenance Intermediate Point (MIP) but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR.

LBM message

The LBM packet is often compared to a ping. A MEP transmits the LBM packet. This packet can be addressed to another MEP or to the MAC address of the MP; in the case of SPBM, this is the SPBM system ID or its virtual SMLT MAC. Only the MP for which the packet is addressed responds with an LBR message.

- Provides “ICMP ping like” functionality natively at Layer-2.
- DA is the MAC address of the target.
- Includes a transaction identifier that allows the corresponding LBR to be identified when more than one LBM request is waiting for a response.
- Bridges forward the frame using the normal FDB rules.
- Only the target (MIP or MEP) responds.
- Initiator can choose the size and contents data portion of the LBM frame.
- Can be used to check the ability of the network to forward different sized frames.

I2 ping

The `l2 ping` command is a proprietary command that allows a user to trigger an LBM message.

- For B-VLANs, specify either the destination MAC address or node name.
- For C-VLANs, specify the destination MAC address.

*** Note:**

To use Layer 2 ping for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the `cfm cmac enable` command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the `l2 ping` command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the `l2 ping` command to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to `l2 ping` and `l2 traceroute` requests.

The `l2 ping` command provides a simpler command syntax than the standard LBM commands, which require you to specify the MD, MA, and MEP ID information. The `l2 ping` command provides a ping equivalent at Layer 2 for use with nodes on the SPBM B-VLAN or C-VLANs.

The options supported for the `12 ping` command vary based on the VLAN type. Only SPBM B-VLANs support the SMLT virtual option for the source mode. Only C-VLANs support the no VLAN MAC option on the source mode.

12 ping with IP address

The `12 ping` command also allows you to specify an IP address as the destination address. In this case, the IP address can be either C-VLAN or B-VLAN in the SPBM cloud.

The `12 ping` command converts Layer 3 IP information to an appropriate Layer 2 VLAN and MAC combination. The system can also target IP addresses that are not SPBM derived routes.

When the `12 ping` command is executed with an IP address as the destination, the operation finds all the valid MAC combinations that provide valid paths to the destination. If ECMP is enabled, there can be multiple paths to the destination. In this case, `12 ping` runs internally for each of the VLAN paths returned, and displays a summary of the results. If ECMP is disabled, the results display only one path.

Fault isolation

Fault isolation is achieved through the use of Linktrace Messages (LTM). LTM is intercepted by all the MPs on the way to the destination MP. Virtual Services Platform 9000 supports two types of LTM.

The first type, the unicast LTM, can be addressed to either MEP or MIP MAC address. Each MP on the way decrements the TTL field in the LTM frame, sends Linktrace Reply (LTR), and forwards the original LTM to the destination. The LTM is forwarded until it reaches its destination or the TTL value is decremented to zero. LTR is a unicast message addressed to the originating MEP.

The second type, the proprietary LTM, is used to map the MAC addresses of the SPBM network; in this case the target MAC is not an MP, but rather a service instance identifier (I-SID).

Link trace message

Connectivity Fault Management offers link trace messaging for fast fault detection. Link trace messages allow operators, service providers and customers to verify the connectivity that they provide or use and to debug systems.

Link trace message — unicast

The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP, which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

- Trace the path to any certain MAC address.
- DA is unicast
- LTM contains:
 - Time to live (TTL)
 - Transaction Identifier

- Originator MAC address
- Target MAC address
- CFM unaware entities forward the frame as is like any other data frame.
- MIP or MEP that is not on the path to the target discards the LTM and does not reply.
- MIP that is on the path to the target:
 - Forwards the LTM after decrementing the TTL and replacing the SA with its own address.
 - Sends a reply (LTR) to the originator.
 - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- If the MIP or MEP is a target:
 - Sends an LTR to the originator.
 - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- A MEP that is not the target but is on the path to the target:
 - Generates a reply as described in the preceding information.
 - It also sets one of the flags fields in the reply to indicate that it is the terminal MEP.

Link trace message — multicast

The multicast link trace message (LTM) can be used to trace the multicast tree from any node on any I- SID using the nickname MAC address and the I-SID multicast address.

Specifying a multicast target address for an LTM allows for the tracing of the multicast tree corresponding to that destination address (DA). With a multicast target every node that is in the active topology for that multicast address responds with a Linktrace reply and also forwards the LTM frame along the multicast path. Missing Linktrace replies (LTRs) from the nodes in the path indicate the point of first failure.

This functionality allows you to better troubleshoot I-SID multicast paths in a SPBM network.

I2 traceroute

The `12 traceroute` command is a proprietary command that allows a user to trigger an LTM message.

- For B-VLANs, specify either the destination MAC address or node name.
- For C-VLANs, specify the destination MAC address.

Note:

To use Layer 2 traceroute for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the `cfm cmac enable` command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the `12 traceroute` command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the `12 traceroute` command to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to `12 ping` and `12 traceroute` requests.

This command provides a simpler command syntax than the standard LTM commands, which require you to specify the MD, MA, and MEP ID information. The `12 traceroute` command provides a trace equivalent at Layer 2 for use with nodes on the SPBM B-VLAN or C-VLANs.

The options supported for the `12 traceroute` command vary based on the VLAN type. Only SPBM B-VLANs support the SMLT virtual option for the source mode. Only C-VLANs support the no VLAN MAC option on the source mode.

12 traceroute with IP address

The `12 traceroute` command allows you to specify an IP address as the destination address. In this case, the IP address can be either C-VLAN or B-VLAN in the SPBM cloud.

The `12 traceroute` command converts Layer 3 IP information to an appropriate Layer 2 VLAN and MAC combination. The system can also target IP addresses that are not SPBM derived routes.

If ECMP is enabled, `12 traceroute` runs internally for each of the VLAN paths returned, and displays a summary of the results. If ECMP is disabled, the results display only one path.

Destination addresses for C-VLAN 12 traceroute and linktrace messages

For C-VLANs, CFM uses the following destination MAC addresses for the corresponding maintenance domain (MD) levels for 12 traceroute and linktrace messages.

VSP 9000 supports both 12 traceroute and linktrace for C-VLANs, but Avaya prefers you use 12 traceroute.

Table 32: MD levels and corresponding destination addresses for CFM for C-VLANs

CFM MD level	Destination MAC address
0	01:80:c2:00:00:38
1	01:80:c2:00:00:39
2	01:80:c2:00:00:3a
3	01:80:c2:00:00:3b
4	01:80:c2:00:00:3c
5	01:80:c2:00:00:3d
6	01:80:c2:00:00:3e
7	01:80:c2:00:00:3f

I2 tracetree

The `I2 tracetree` command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

VSP 9000 only supports this command on SPBM B-VLANs. VSP 9000 does not support `I2 tracetree` command for C-VLANs.

Layer 2 tracemroute

The `I2tracemroute` command is a proprietary command that allows the user to trace the multicast tree for a certain multicast flow. The user specifies source, group, and service context (either VLAN or VRF) for the multicast flow to trace.

CFM sends a multicast LTM using an internal calculation to map the source, group, and context to the corresponding target address. The LTR comes from all leaves of the multicast tree for that flow, as well as transit nodes. The target MAC used in the LTM is a combination of the data I-SID and the nickname and the packet is sent on the appropriate SPBM B-VLAN. The user can see the generated multicast tree for that flow, which includes the data I-SID and nickname.

Nodal MPs

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. The Nodal MEP provides traceability and troubleshooting at the system level for a certain B-VLAN. Each node (chassis) has a certain MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP. The Nodal B-VLAN MPs supports eight levels of CFM and you configure the Nodal B-VLAN MPs on a per B-VLAN basis. Virtual SMLT MAC addresses are also able to respond for LTM and LBM.

Nodal B-VLAN MEPs

The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs. To support this behavior a MAC Entry is added to the FDB and a new CFM data-path table containing the B-VLAN and MP level are added to direct CFM frames to the CP as required.

Nodal B-VLAN MIPs

The Nodal MIP is associated with a B-VLAN. VLAN and level are sufficient to specify the Nodal MIP entity. The Nodal MIP MAC address is the SPBM system ID for the node on which it resides. If the

fastpath sends a message to the CP, the MIP responds if it is not the target and the MEP responds if it is the target.

Nodal B-VLAN MEPs and MIPs with SMLT

When Nodal MEPs or MIPs are on SPBM B-VLANs the LTM code uses a unicast MAC DA. The LTM DA is the same as the target MAC address, which is the SPBM MAC address or the SMLT MAC address of the target node.

Virtual Services Platform 9000 supports SMLT interaction with SPBM. This is accomplished by using two B-VIDs into the core from each pair of SMLT terminating nodes. Both nodes advertise the Nodal B-MAC into the core on both B-VIDS. In addition each node advertises the SMLT virtual B-MAC on one of the two B-VLANs.

The Nodal MEP and MIP are expanded to respond to both the Nodal MAC address as well as the Virtual SMLT MAC address if both MACs are being advertised on its B-VLAN. In addition a source mode is added to the LTM and LBM command to use either the Nodal MAC or the SMLT virtual MAC address as the source MAC in the packet.

Configuration considerations

When you configure CFM, be aware of the following configuration considerations:

- A single switch has a limit of two nodal MEPs and two nodal MIPs
- All nodal MEPs and MIPs are restricted to SPBM B-VIDs.
- The Maintenance level for MEPs and MIPs on a certain B-VID (in a network) must be configured to the same level for them to respond to a certain CFM command.

Limitations

When you configure CFM, be aware of the following configuration limitations:

- CFM does not support CCM messages.
- Only an autogenerated MEP and MIP can exist for C-VLANs.
- Only one MEP can exist on each C-VLAN and SPBM B-VLAN.
- Only one MIP can exist on each C-VLAN and SPBM B-VLAN.
- SMLT Virtual MAC for C-VLAN does not exist, so VSP 9000 does not support this option for I2 ping and I2 traceroute.
- VSP 9000 does not support I2tracetree on C-VLANs because no multicast tree exists on C-VLANs.
- The autogenerated MEPs do not have a uniqueness across the entire network until you configure the global MEP ID on each box to a different value. You must configure a unique MEP ID at a global level for CFM.
- MEPs and MIPs configured for SPBM VLANs do not respond to CFM messages sent from C-MAC VLANs because the VLAN and packet encapsulation are different.
- Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.

- You can only configure global CFM at one MD level for each chassis for each VLAN type.

CFM configuration using ACLI

This section provides procedures to configure and use Connectivity Fault Management (CFM) using Avaya Command Line Interface (ACLI). The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults, which is performed at Layer 2, not Layer 3. To support troubleshooting of the SPBM cloud, Virtual Services Platform 9000 supports a subset of CFM functionality.

*** Note:**

When you enable CFM in an SPBM network, Avaya recommends that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

Autogenerated CFM

CFM provides two methods for configuration: autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure a MD, MA, and MEP ID to create a MEP.

*** Note:**

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.

VSP 9000 only supports one MEP and one MIP, either autogenerated or explicitly configured, on the SPBM B-VLAN. Similarly, VSP 9000 only supports one MEP and one MIP on the C-VLAN. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN.

For autogenerated CFM configuration information for ACLI see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 197.
- [Configuring autogenerated CFM on C-VLANs](#) on page 198.

Configuring autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This eliminates the need to explicitly configure an MD, MA, and MEP ID, and to associate the MEP and MIP level to the SPBM B-VLAN.

When you enable this feature, the device creates a global MD (named `spbm`) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the `level` attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1.

The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

Important:

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs:

```
cfm spbm level <0-7>
```

You can change this level from the default of 4 either before or after the feature is enabled.

Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM SPBM MEPs:

```
cfm spbm mepid <1-8191>
```

4. Enable the autogenerated CFM for SPBM B-VLANs globally:

```
cfm spbm enable
```

5. **(Optional)** Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs to the default:

```
default cfm spbm level
```

6. **(Optional)** Assign a global CFM MEP ID for all CFM SPBM MEPs to the default:

```
default cfm spbm mepid
```

7. **(Optional)** Disable the global CFM MEPs and MIPs:

```
no cfm spbm enable
```

8. Display the global CFM MEP configuration:

```
show cfm spbm
```

Example

Configure autogenerated CFM MEPs and MIPs:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#cfm spbm level 6
VSP-9012:1(config)#cfm spbm mepid 4
VSP-9012:1(config)#cfm spbm enable
VSP-9012:1(config)#show cfm spbm

LEVEL ADMIN      MEPID      MAC
=====
6          enable          4          00:15:e8:b8:a3:df
```

Variable definitions

Use the data in the following table to use the `cfm spbm` command.

Variable	Value
level<0-7>	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
mepid<1-8191>	Specifies the global MEP ID within the range of 1–8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
enable	Enables autogenerated CFM on all SPBM B-VLANs.

Job aid

The following table describes the fields for the `show cfm spbm` command.

Parameter	Description
LEVEL	Specifies the global SPBM CFM maintenance level for the chassis. The default is 4.
ADMIN	Specifies if CFM MEPs and MIPs are globally enabled.
MEP ID	Specifies the global MEP ID. The default is 1.
MAC	Specifies the MAC address.

Configuring autogenerated CFM on C-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every C-VLAN on the chassis.

! Important:

CFM supports one MEP or MIP for each C-VLAN. Only autogenerated CFM provides support for configuring MEP and MIPs on C-VLANs. You cannot explicitly configure C-VLANs.

About this task

When you enable this feature, you create a global MD (named `cmac`) for all the customer MAC (C-MAC) MEPs. This global MD has a default maintenance level of 4, which you can change with the `level` attribute. The autogenerated CFM commands also create an MA for each C-VLAN, a MEP for each C-VLAN, associate the MEP with the corresponding C-VLAN, and a MIP with the C-VLAN.

All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The device automatically associates the MEPs with the C-VLANs configured. The MIP level maps to the global level. The device automatically associates the MIP level with the C-VLANs configured when you enable the feature.

The feature is disabled by default.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the maintenance level for every CFM C-MAC MEP and MP level on all the C-VLANs:

```
cfm cmac level <0-7>
```

Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM C-MAC MEPs:

```
cfm cmac mepid <1-8191>
```

4. Enable the autogenerated CFM for C-VLANs:

```
cfm cmac enable
```

5. **(Optional)** Configure the maintenance level for every CFM C-MAC MEPs and MP level on all the C-VLANs to the default:

```
default cfm cmac level
```

6. **(Optional)** Assign a global CFM MEP ID for all CFM C-MAC MEPs to the default:

```
default cfm cmac mepid
```

7. **(Optional)** Disable the global CFM MEPs and MIPs:

```
no cfm cmac enable
```

8. Display the global CFM MEP configuration:

```
show cfm cmac
```

Example

Configure autogenerated CFM MEPs and MIP level:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#cfm cmac level 0
VSP-9012:1(config)#cfm cmac mepid 4
VSP-9012:1(config)#cfm cmac enable
VSP-9012:1(config)#show cfm cmac

LEVEL ADMIN      MEPID      MAC
=====
0          enable          4          00:15:e8:b8:a3:de
```

Variable definitions

Use the data in the following table for the `cfm cmac` command.

Variable	Value
level<0-7>	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
mepid<1-8191>	Specifies the global MEP ID within the range of 1–8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1. * Note: The MA takes its name from this value for autogenerated CFM. For example, if you specify 500 as the MEP ID, the MA will also be 500.
enable	Enables autogenerated CFM for all C-MAC VLANs.

Job aid

The following table describes the fields for the `show cfm cmac` command.

Parameter	Description
LEVEL	Specifies the global C-VLAN CFM maintenance level for the chassis. The default is 4.
ADMIN	Specifies if CFM C-VLAN MEPs and MIPs are globally enabled.
MEP ID	Specifies the global MEP ID. The default is 1.
MAC	Specifies the MAC address.

Configuring explicit CFM

For SPBM B-VLANs, CFM provides two methods for configuration: autogenerated and explicit. You cannot use both. Use the procedures in this section to configure CFM explicitly. For C-VLANs, you can only use the autogenerated method.

If you want to create explicit CFM MEPs that require you to configure an MD, MA, and MEP ID, see the procedures in the following sections:

- [Configuring CFM MD](#) on page 201.
- [Configuring CFM MA](#) on page 202.
- [Assigning a MEP and MIP level to an SPBM B-VLAN](#) on page 203.
- [Configuring CFM MEP](#) on page 205.

* Note:

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

Configuring CFM MD

Use this procedure to configure the Connectivity Fault Management (CFM) Maintenance Domain (MD) explicitly. An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

* Note:

If you use autogenerated CFM, you do not configure CFM MD because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create the CFM MD:

```
cfm maintenance-domain WORD<1-22> [index <1-2147483647>]
[maintenance-level <0-7>] [level <0-7>]
```

3. Display the CFM MD configuration:

```
show cfm maintenance-domain
```

4. Delete the CFM MD:

```
no cfm maintenance-domain WORD<1-22>
```

Example

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

```
VSP-9012:1# configure terminal
VSP-9012:1(config)# cfm maintenance-domain mdl index 99 maintenance-level
3
VSP-9012:1(config)# show cfm maintenance-domain
```

```
=====
Maintenance Domain
=====
Domain Name          Domain Index    Level Domain Type
-----
mdl                  99              3      NONE
Total number of Maintenance Domain entries: 1.
```

```
VSP-9012:1(config)# no cfm maintenance-domain mdl
VSP-9012:1(config)#show cfm maintenance-domain
```

```
=====
Maintenance Domain
=====
Domain Name          Domain Index    Level Domain Type
-----
Total number of Maintenance Domain entries: 0.
```

Variable definitions

Use the data in the following table to use the **cfm maintenance-domain** command.

Variable	Value
WORD<1-22>	Specifies the maintenance domain name.
index <1-2147483647>	Specifies a maintenance domain entry index.
maintenance-level <0-7>	Specifies the MD maintenance level when creating the MD. The default is 4.
level <0-7>	Modifies the MD maintenance level for an existing MD. The default is 4.

Configuring CFM MA

Use this procedure to configure the CFM Maintenance Association (MA) explicitly. An MA represents a logical grouping of monitored entities within its domain. It can therefore represent a set of Maintenance Association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

*** Note:**

If you use autogenerated CFM, you do not configure CFM MA because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. Enter Global Configuration mode:
enable

```
configure terminal
```

2. Create the CFM MA:

```
cfm maintenance-association WORD<1-22> WORD<1-22> [index <1-2147483647>]
```

3. Display the CFM MA configuration:

```
show cfm maintenance-association
```

4. Use the following command, if you want to delete the CFM MA:

```
no cfm maintenance-association WORD<1-22> WORD<1-22>
```

Example

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

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

```
VSP-9012:1(config)# cfm maintenance-association mdl mal index 98
```

```
VSP-9012:1(config)# show cfm maintenance-association
```

```
=====
Maintenance Association Status
=====
Domain Name          Assn Name          Domain Idx  Assn Idx
-----
mdl                  mal                1           98
Total number of Maintenance Association entries: 1.
=====
Maintenance Association config
=====
Domain Name          Assn Name
-----
mdl                  mal
Total number of MA entries: 1.
```

Variable definitions

Use the data in the following table to use the **cfm maintenance-association** command.

Variable	Value
WORD<1-22> WORD<1-22>	Creates the CFM MA. The first parameter, specifies the MD name. The second parameter, specifies the MA name.
index <1-2147483647>	Specifies a maintenance association entry index.

Assigning a MEP and MIP level to an SPBM B-VLAN

Use this procedure to assign a nodal MEP to an SPBM B-VLAN. The Nodal MEP provides traceability and troubleshooting at the system level for a specific B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the specific B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a specific MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

*** Note:**

If you use autogenerated CFM, you do not configure CFM MIP/MEP because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Add nodal MEPs to the B-VLAN:

```
vlan nodal-mep <1-4084> WORD<0-22> WORD<0-22> <1-8191>
```

3. Display the nodal MEP configuration:

```
show vlan nodal-mep <1-4084>
```

4. Remove the nodal MEPs from the B-VLAN:

```
no vlan nodal-mep <1-4084> WORD<0-22> WORD<0-22> <1-8191>
```

5. Add nodal MIP level to the B-VLAN:

```
vlan nodal-mip-level <1-4084> WORD<0-15>
```

6. Display the nodal MIP level configuration:

```
show vlan nodal-mip-level [<1-4084>]
```

7. Remove the nodal MIP level from the B-VLAN:

```
no vlan nodal-mip-level <1-4084> WORD<0-15>
```

Example

```
VSP-9012:1> enable
VSP-9012:1> configure terminal
VSP-9012:1> vlan nodal-mep 100 md1 ma1 2
VSP-9012:1> show vlan nodal-mep
```

```
=====
                        Vlan Nodal Mep
=====
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
-----
1
```

```

100      md1.ma1.2
216
304
404
500
616
716
816
916
1000    spbm.1000.1
1001    spbm.1001.1

```

```
VSP-9012:> vlan nodal-mip 100 6
```

```
VSP-9012:> show vlan nodal-mip
```

```

=====
                        Vlan Nodal Mip Level
=====
VLAN_ID   NODAL_MIP_LEVEL_LIST
-----
1
100       6
216
304
404
500
616
716
816
916
1000     4
1001     4

```

Variable definitions

Use the data in the following table to use the `vlan nodal-mep` command.

Variable	Value
<1-4084>	Specifies the VLAN ID.
WORD<0-22>	The first parameter, specifies the Maintenance Domain name.
WORD<0-22>	The second parameter, specifies the Maintenance Association name.
<1-8191>	Specifies the nodal MEPs to add to the VLAN.

Use the data in the following table to use the `vlan nodal-mip-level` command.

Variable	Value
<1-4084>	Adds the nodal MIP level. Specifies the VLAN ID.
WORD<0-15>	Adds the nodal MIP level, which has up to eight levels, ranging from 0 to 7.

Configuring CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service

instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

*** Note:**

If you use autogenerated CFM, you do not configure CFM MEPs because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create the CFM MEP:

```
cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191> [state
<enable>]
```

3. Enable an existing CFM MEP:

```
cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191> enable
```

4. Disable an existing CFM MEP:

```
no cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191> enable
```

5. Display the CFM MEP configuration:

```
show cfm maintenance-endpoint
```

6. Delete an existing CFM MEP:

```
no cfm maintenance-endpoint WORD<1-22> WORD<1-22> <1-8191>
```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# cfm maintenance-endpoint md1 ma1 1 state enable
VSP-9012:1> show cfm maintenance-endpoint
```

```
=====
Maintenance Endpoint Config
=====
```

DOMAIN NAME	ASSOCIATION NAME	MEP ID	ADMIN
md1	ma1	1	enable

```
-----
Total number of MEP entries: 1.
=====
Maintenance Endpoint Service
=====
```

DOMAIN_NAME	ASSN_NAME	MEP_ID	TYPE	SERVICE_DESCRIPTION
-------------	-----------	--------	------	---------------------

```
-----
```

```
md1          ma1          1          unused
Total number of MEP entries: 1.
```

Variable definitions

Use the data in the following table to use the `cfm maintenance-endpoint` command.

Variable	Value
<code>WORD<1-22></code>	The first parameter, specifies the MD name.
<code>WORD<1-22></code>	The second parameter, specifies the MA name.
<code><1-8191></code>	Specifies the MEP ID.
<code>state {enable disable}</code>	Enables or disables the MEP when creating the MEP. The default is disabled.
<code>enable</code>	Enables an existing MEP. Use this parameter with the <code>no</code> option to disable an existing MEP.

Triggering a loopback test (LBM)

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before you begin

- You must have a MEP that is associated with a B-VLAN or a C-VLAN.

Procedure

- Log on to the switch to enter User EXEC mode.
- Trigger the loopback test:

```
loopback WORD<1-22> WORD<1-22> <1-8191>
<0x00:0x00:0x00:0x00:0x00:0x00> [burst-count <1-200>] [data-tlv-size
<0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode
{nodal|noVlanMac|smltVirtual}] [testfill-pattern <all-zero|all-zero-
crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>]
[time-out <1-10>]
```

Example

```
VSP-9012:1# loopback md1 4001 13 00:14:0D:A2:B3:DF burst-count 10 priority
3 time-out 5
```

```
Result of LBM from mep: md1.4002.13 to MAC address: 00:14:0D:A2:B3:DF :
Sequence number of the first LBM is 10575
The total number of LBMs sent out is 10
The number of LBRs received is 10
The number of LBRs lost is 0
The percentage of LBMs lost is 0.00%
```

The RTT Min is 764 microsecs, Max is 800 microsecs, Average is 783.00 microsecs
 The RTT DV min is 3 microsecs, Max is 23 microsecs, Average is 9.11 microsecs
 The Standard Deviation of RTT is 11.53 microsecs

Variable definitions

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

Variable	Value
<code>WORD<1–22></code>	The first parameter, specifies the MD name.
<code>WORD<1–22></code>	The second parameter, specifies the MA name.
<code><1–8191></code>	Specifies the MEP ID.
<code><0x00:0x00:0x00:0x00:0x00:0x00></code>	Specifies the remote MAC address to reach the MEP/MIP.
<code>burst-count <1–200></code>	Specifies the burst-count.
<code>data-tlv-size <0–400></code>	Specifies the data TLV size.
<code>frame-size <64–1500></code>	Specifies the frame-size. The default is 0.
<code>priority <0–7></code>	Specifies the priority. The default is 7.
<code>source-mode {nodal noVlanMac smltVirtual}</code>	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the B-MAC-SA. • smltVirtual—Use this value with B-VLANs only. <p>The default is nodal.</p>
<code>testfill-pattern {all-zero all-zero-crc pseudo-random-bit-sequence pseudo-random-bit-sequence-crc}</code>	<p>Specifies the testfill pattern:</p> <ul style="list-style-type: none"> • all-zero — null signal without cyclic redundancy check • all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial • pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors.</p> <p>The default is 1: all-zero.</p>
<code>time-out <1–10></code>	Specifies the time-out interval in seconds. The default is 3.

Triggering linktrace (LTM)

Use the following procedure to trigger a linktrace.

The Linktrace Message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Before you begin

- You must have a MEP that is associated with a B-VLAN or C-VLAN.

Procedure

- Log on to the switch to enter User EXEC mode.
- Trigger the linktrace:

```
linktrace WORD<1-22> WORD<1-22> <1-8191>
<0x00:0x00:0x00:0x00:0x00:0x00> [detail] [priority <0-7>] [source-
mode <nodal|noVlanMac|smltVirtual>] [ttl-value <1-255>]
```

Example

```
VSP-9012:1# linktrace mdl 4001 13 00:bb:00:00:14:00 priority 7
```

```
Please wait for LTM to complete or press any key to abort
```

```
Received LTRs:
```

```
SeqNum: 10575 MD: mdl MA:4001 MepId: 13 Priority: 7
-----
TTL SRC MAC FWDYES TERMMEP RELAY ACTION
-----
63 00:bb:00:00:10:00 true false Fdb
62 00:bb:00:00:14:00 false true Hit
```

Variable definitions

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

Variable	Value
<i>WORD</i> <1-22>	The first parameter, specifies the MD name.
<i>WORD</i> <1-22>	The second parameter, specifies the MA name.
<1-8191>	Specifies the MEP ID.
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the target MAC address to reach the MEP.
detail	Displays linktrace result details.
priority <0-7>	Specifies the priority. The default is 7.
source-mode <nodal noVlanMac smltVirtual>	Specifies the source mode: <ul style="list-style-type: none"> 1: nodal noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. 2: smltVirtual—Use this value with B-VLANs only.

Table continues...

Variable	Value
	The default is 1: nodal.
ttl-value <1-255>	Specifies the Time-to-Live value. The default is 64.

Triggering a Layer 2 ping

Use this procedure for C-VLANs or B-VLANs to trigger a Layer 2 ping, which acts like native `ping`. This feature enables CFM to debug Layer 2. Layer 2 ping can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

* Note:

To use Layer 2 ping for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the `cfm cmac enable` command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the `l2 ping` command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the `l2 ping` command to test reachability for all the B-MAC addresses in the SPBM network.

* Note:

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to `l2 ping` and `l2 traceroute` requests.

Before you begin

- You must configure and enable CFM.
- You must have a MEP that is associated with a VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 ping:

```
l2 ping {vlan <1-4084>} {routernodename WORD<0-255> | mac
<0x00:0x00:0x00:0x00:0x00:0x00>} [burst-count <1-200>] [data-tlv-
size <0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode
<nodal|noVlanMac|smltVirtual>] [testfill-pattern <all-zero|all-zero-
crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>]
[time-out <1-10>]
```

```
l2 ping {ip-address WORD<0-255>} [burst-count <1-200>] [data-tlv-
size <0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode
<nodal|noVlanMac|smltVirtual>] [testfill-pattern <all-zero|all-zero-
```

```
crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>]
[time-out <1-10>] [vrf WORD<1-16>]
```

Example

Trigger a Layer 2 ping for MAC address 00.14.0d.bf.a3.d:

```
VSP-9012:1> l2 ping vlan 500 mac 00.14.0d.bf.a3.d
Please wait for l2ping to complete or press any key to abort
----00:14:0d:bf:a3:df L2 PING Statistics---- 0(64) bytes of data
1 packets transmitted, 0 packets received, 100.00% packet loss
```

Trigger a Layer 2 ping for router node name VSP-MONTIO:

```
VSP-9012:1> l2 ping vlan 500 routernodename VSP-MONTIO
Please wait for l2ping to complete or press any key to abort

----00:14:0d:a2:b3:df L2 PING Statistics---- 0(64) bytes of data
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip (us) min/max/ave/stdv = 26895/26895/26895.00/ 0.00
```

Trigger a Layer 2 ping for IP address 192.0.2.102:

```
VSP-9012:1>l2 ping ip-address 192.0.2.102

Please wait for l2ping to complete or press any key to abort

L2 PING Statistics : IP 192.0.2.102, paths found 1, path attempted 1
=====
VLAN NEXT HOP TX RX PERCENT ROUND TRIP TIME
PKTS PKTS LOSS MIN/MAX/AVE (us)
=====
50 80:17:7d:75:aa:02 (80:17:7d:75:aa:02) 1 0 100.00% 0/0/0.00
```

Variable definitions

Use the data in the following table to configure the `l2 ping` command.


Variable	Value
{vlan <1-4084> routernodename WORD<0-255>} (vlan <1-4084> mac <0x00:0x00:0x00:0x00:0x00:0x00>) {ip-address WORD<0-255>}	Specifies the destination for the L2 ping: <ul style="list-style-type: none"> <1-4084> — Specifies the VLAN ID. WORD<0-255> — Specifies the Router node name. <XX:XX:XX:XX:XX:XX> — Specifies the MAC address. <A.B.C.D> — Specifies the IP address. <p> Note: VSP 9000 does not support the routernodename option for C-VLANs.</p>
burst-count <1-200>	Specifies the burst count.
data-tlv-size <0-400>	Specifies the data TLV size. The default is 0.
frame-size <64-1500>]	Specifies the frame size. The default is 0.

Table continues...

Variable	Value
testfill-pattern <all-zero all-zero-crc pseudo-random-bit-sequence pseudo-random-bit-sequence-crc>	<p>Specifies the testfill pattern:</p> <ul style="list-style-type: none"> all-zero — Null signal without cyclic redundancy check. all-zero-crc — Null signal with cyclic redundancy check with 32-bit polynomial. pseudo-random-bit-sequence — Pseudo-random-bit-sequence without cyclic redundancy check. pseudo-random-bit-sequence-crc — Pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors. The default is all-zero.</p>
priority <0–7>]	Specifies the priority. The default is 7.
time-out <1–10>	Specifies the interval in seconds. The default is 3.
source-mode <nodal noVlanMac smltVirtual>	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> nodal noVlanMac — Use this value with C-VLANs only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the B-MAC-SA. smltVirtual—Use this value with B-VLANs only. <p>The default is nodal.</p>
vrf WORD<1–16>	Specifies the VRF name.

Triggering a Layer 2 traceroute

Use this procedure for B-VLANs or C-VLANs to trigger a Layer 2 traceroute, which acts like native **traceroute**. This feature enables CFM to debug Layer 2 for SPBM B-VLANs or C-VLANs. Layer 2 traceroute can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

Note:

To use Layer 2 traceroute for C-MAC addresses to test connectivity between access points, you must configure CFM C-MAC with the **cfm cmac enable** command and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with the **l2 traceroute** command, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **l2 traceroute** command to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to `12 ping` and `12 traceroute` requests.

Before you begin

- You must configure and enable CFM.
- You must have a MEP that is associated with a VLAN.

About this task

! Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an `12 ping` to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.
- For C-VLANs, you have to trigger an `12 ping` to learn the C-MAC address.

In both cases, `linktrace` traces the path up to the closest device to that MAC address that supports CFM.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 traceroute:

```
12 traceroute {vlan <1-4084> routernodename WORD<0-255> | vlan <1-4084> mac <0x00:0x00:0x00:0x00:0x00:0x00>} [priority <0-7>] [source-mode <nodal|noVlanMac|smltVirtual>] [ttl-value <1-255>]

12 traceroute {ip-address WORD<0-255>} [priority <0-7>] [source-mode <nodal|noVlanMac|smltVirtual>] [ttl-value <1-255>] [vrf WORD<1-16>]
```

Example

Trigger a Layer 2 traceroute for VLAN 500 with the router node name VSP-MONTIO:

```
VSP-9012:1#12 traceroute vlan 500 routernodename VSP-MONTIO

Please wait for 12traceroute to complete or press any key to abort

12traceroute to VSP-MONTIO (00:14:0d:a2:b3:df), vlan 500
0   VSP-PETER4 (00:15:9b:11:33:df)
1   VSP-MONTIO (00:14:0d:a2:b3:df)
```

Trigger a Layer 2 traceroute for the IP address 192.0.2.1:

```
VSP-9012:1#12 traceroute ip-address 192.0.2.1

Please wait for 12trace to complete or press any key to abort

L2 Trace Statistics : IP 192.0.2.1, paths found 1
=====
```

```
VSP-SHAMIM (00:1a:8f:08:53:df), vlan 500
0 VSP-PETER4 (00:15:9b:11:33:df)
1 VSP-MONTI0 (00:14:0d:a2:b3:df)
```

The output for the L2 traceroute using C-VLAN 10 to target MAC 00:14:9b:11:30:00.

```
VSP-9012:1# l2 traceroute 10.00:14:9b:11:30:00

Please wait for l2traceroute to complete or press any key to abort
l2traceroute to 00:14:9b:11:30:00, vlan 10
0 00:15:9b:11:30:00 (00:15:9b:11:30:00)
1 00:14:9b:11:30:00 (00:14:9b:11:30:00)
```

Variable definitions

Use the data in the following table to use the `l2 traceroute` command.

Variable	Value
{vlan <1–4084> routernodename WORD<0–255>} (vlan <1–4084> mac <0x00:0x00:0x00:0x00:0x00:0x00>) {ip-address WORD<0–255>}	Specifies the destination for the L2 traceroute: <ul style="list-style-type: none"> • <1–4084> — Specifies the VLAN ID. • WORD<0–255> — Specifies the router node name. • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address. • WORD<0–255> — Specifies the IP address. <p>* Note: VSP 9000 does not support the routernodename option for C-VLANs.</p>
ttl-value<1–255>	Specifies the TTL value. The default is 64.
priority <0–7>	Specifies the priority. The default is 7.
source-mode <nodal noVlanMac smltVirtual>	Specifies the source mode: <ul style="list-style-type: none"> • nodal • noVlanMac — Use this value with C-VLANs only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the B-MAC-SA. • smltVirtual— Use this value with B-VLANs only. The default is nodal.
vrf WORD<1–16>	Specifies the VRF name.

Triggering a Layer 2 tracetree

Use this procedure to trigger a Layer 2 tracetree. Layer 2 tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

*** Note:**

VSP 9000 only supports this command on SPBM B-VLANs only, not C-VLANs.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 tracetable:

```
12 tracetable {<1-4084> <1-16777215> [routernodename WORD<0-255> |
<1-4084> <1-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>]}
[priority <0-7>] [source-mode <nodal|smltVirtual>] [ttl-value <1-
255>]
```

Example

```
VSP-9012:1# 12 tracetable 500 1
```

```
Please wait for 12tracetable to complete or press any key to abort
```

```
12tracetable to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10
hops 64
1 VSP-PETER4 00:15:9b:11:33:df -> VSP-MONTI0 00:14:0d:a2:b3:df
2 VSP-MONTI0 00:14:0d:a2:b3:df -> VSP-LEE2 00:15:e8:b8:a3:df
```

Variable definitions

Use the data in the following table to use the `12 tracetable` command.

Variable	Value
{<1-4084> <1-16777215> routernodename WORD<0-255> <1-4084> <1-16777215> mac <0x00:0x00:0x00:0x00:0x00:0x00>}	<ul style="list-style-type: none"> • <1-4084> — Specifies the VLAN ID. • <1-16777215> — Specifies the I-SID. • WORD<0-255> — Specifies the Router Node Name. • <0x00:0x00:0x00:0x00:0x00:0x00> — Specifies the MAC address.
ttl-value<1-255>	Specifies the TTL value. The default is 64.
priority <0-7>	Specifies the priority value. The default is 7.
source-mode <nodal smltVirtual>	Specifies the source mode. <ul style="list-style-type: none"> • 1: nodal • 2: smltVirtual The default is nodal.

Triggering a Layer 2 tracemroute

Use this procedure to debug the IP multicast over SPBM stream path using `l2 tracemroute` on the VLAN (Layer 2) or the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

* Note:

The VLAN option is only valid for a VLAN that has an I-SID configured and IGMP snooping enabled.

Before you begin

- On the source and destination nodes, you must configure an autogenerated or an explicit CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Trigger a Layer 2 tracemroute on the VLAN:

```
l2 tracemroute source <A.B.C.D> group <A.B.C.D> vlan
<1-4084>[priority <0-7>] [ttl-value <1-255>]
```

* Note:

For the preceding command, if you do not specify a VLAN, `l2 tracemroute` uses the global default VRF.

Wait for the l2 tracemroute to complete or press any key to abort.

3. Trigger a Layer 2 tracemroute on the VRF:

```
l2 tracemroute source <A.B.C.D> group <A.B.C.D> vrf WORD<1-16>
[priority <0-7>] [ttl-value <1-255>]
```

* Note:

For the preceding command, if you do not specify a VRF, `l2 tracemroute` uses the global default VRF.

Wait for the l2 tracemroute to complete or press any key to abort.

Example

The following is a sample output for a Layer 2 tracemroute on a VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#l2 tracemroute source 192.0.2.81 233.252.0.1 vlan 201
Please wait for l2 tracemroute to complete or press any key to abort.
```



```

Source 192.0.2.81
Group: 233.252.0.1
VLAN:201
BMAC: 03:00:03:f4:24:01
B-VLAN: 10
I-SID: 16000001

=====
1 VSP-PETER4 00:03:00:00:00:00 -> VSP-LEE1 00:14:0d:bf:a3:df
2 VSP-LEE1 00:14:0d:bf:a3:df -> VSP-LEE2 00:15:e8:b8:a3:df

```

The following is a sample output for a Layer 2 tracemroute on a VRF:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#12 tracemroute source 192.0.2.10 group 233.252.0.1 vrf red

Please wait for 12 tracemroute to complete or press any key to abort.

Source 192.0.2.10
Group: 233.252.0.1
VRF: redID 1
BMAC: 03:00:04:f4:24:01
B-VLAN: 20
I-SID: 16000001

=====
1 VSP-PETER4 00:03:00:00:00:00 -> VSP-LEE1 00:14:0d:bf:a3:df
2 VSP-LEE1 00:14:0d:bf:a3:df -> VSP-LEE2 00:15:e8:b8:a3:df

```

Variable definitions

Use the data in the following table to use the `12 tracemroute` command.

Variable	Value
source <A.B.C.D>	Specifies the source IP address.
group <A.B.C.D>	Specifies the IP address of the multicast group.
vlan <1-4084>	Specifies the VLAN value.
vrf WORD<1-16>	Specifies the VRF name. If you do not specify a VRF name, then the results are shown for the flow in the Global Router (default) context.
priority <0-7>	Specifies the priority value.
ttl <1-255>	Specifies the time-to-live (TTL) for the trace packet, which is how many hops the trace packet takes before it is dropped.

Job aid

The following table describes the fields in the output for 12 `tracemroute` command for a VLAN.

Parameter	Description
Source	Specifies the source IP address of the flow where the multicast trace tree originates.
Group	Specifies the IP address of the multicast group.
VLAN	Specifies the VLAN.
BMAC	Specifies the backbone MAC address.
B-VLAN	Specifies the backbone VLAN.
I-SID	Specifies the service identifier.

The following table describes the fields in the output for 12 `tracemroute` command for a VRF.

Parameter	Description
Source	Specifies the source IP address of the flow where the multicast trace tree originates.
Group	Specifies the IP address of the multicast group.
VRF	Specifies the VRF.
BMAC	Specifies the backbone MAC address.
B-VLAN	Specifies the backbone VLAN.
I-SID	Specifies the service identifier.

Using trace CFM to diagnose problems

Use the following procedure to display trace information for CFM.

About this task

Use trace to observe the status of a software module at a certain time.

For example, if you notice a CPU utilization issue (generally a sustained spike above 90%) perform a trace of the control plane activity.

Use the `trace level 120 <0-4>` command to trace CFM module information, including ACLI, instrumentation, High Availability, show config, and platform dependent code. The CFM module ID is 120.

Use the `trace cfm level <0-4>` command to trace platform independent code and CFM protocol code.

Caution:

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Clear the trace:

```
clear trace
```

3. Begin the trace operation:

```
trace cfm level <0-4>
```

Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

```
trace shutdown
```

5. View the trace results:

```
show trace cfm
```

6. Begin the trace operation for the CFM module:

```
trace level 120 <0-4>
```

Wait approximately 30 seconds, and then stop trace.

7. View trace results:

```
trace screen enable
```

Important:

If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

8. Save the trace file to the Compact Flash card for retrieval.

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

9. Search trace results for a specific string value, for example, the word error:

```
trace grep [WORD<0-128>]
```

If you use this command and do not specify a string value, you clear the results of a previous search.

Example

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

```
VSP-9012:1(config)# clear trace
VSP-9012:1(config)# trace cfm level 3
VSP-9012:1(config)# trace shutdown
VSP-9012:1(config)# show trace cfm
=====
                                CFM Tracing Info
=====
Status      : Enabled
Level       : VERBOSE
VSP-9012:1(config)#trace level 120 3
VSP-9012:1(config)# save trace
VSP-9012:1(config)# trace grep error
VSP-9012:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

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

Table 33: Variable definitions

Variable	Value
cfm level [<i><0-4></i>]	Starts the trace by specifying the level. <ul style="list-style-type: none"> <i><0-4></i> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
grep [<i>WORD<0-128></i>]	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
level <i><0-192></i> [<i><0-4></i>]	Starts the trace by specifying the module ID and level. <ul style="list-style-type: none"> <i><0-192></i> specifies the module ID. <i><0-4></i> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
shutdown	Stops the trace operation.
screen {disable enable}	Enables the display of trace output to the screen.

Use the data in the following table to use the **save trace** command.

Table 34: Variable definitions

Variable	Value
file <i>WORD<1-99></i>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> a.b.c.d:<file> x:x:x:x:x:x:x <file> /intflash/<file> /extflash/<file> /usb/<file> /mnt/intflash/ <file>

Variable	Value
	<ul style="list-style-type: none"> • /mnt/extflash/ <file> <p>/mnt/intflash is the internal flash of the second CP module (the one to which you are not connected).</p> <p>/mnt/extflash is the external flash of the second CP module (the one to which you are not connected).</p>

Using trace SPBM to diagnose problems

Use the following procedure to display trace information for SPBM IS-IS. In the case of IS-IS, this procedure also provides information related to the flags set.

About this task

Use the `trace level 119 <0-4>` command to trace IS-IS module information, including ACLI, instrumentation, High Availability, show config and platform dependent code. The IS-IS module ID is 119.

Use the `trace level 125 <0-4>` command to trace SPBM module information, including ACLI, instrumentation, High Availability, show config and platform dependent code. The SPBM module ID is 125.

Use the `trace spbm isis level` command to trace platform independent code, IS-IS protocol, IS-IS hello, IS-IS adjacency, LSP processing, and IS-IS computation.

Caution:

Risk of traffic loss

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Clear the trace:

```
clear trace
```

3. Begin the trace operation:

```
trace spbm isis level <0-4>
```

Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

```
trace shutdown
```

5. Display the trace information for SPBM IS-IS:

```
show trace spbm isis
```

6. Begin the trace operation for the SPBM module:

```
trace level 125 <0-4>
```

Wait approximately 30 seconds, and then stop trace.

7. Begin the trace operation for the IS-IS module:

```
trace level 119 <0-4>
```

Wait approximately 30 seconds, and then stop trace.

8. View trace results:

```
trace screen enable
```

! **Important:**

If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

9. Save the trace file to the Compact Flash card for retrieval.

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

10. Search trace results for a specific string value, for example, the word error:

```
trace grep [WORD<0-128>]
```

If you use this command and do not specify a string value, you clear the results of a previous search.

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)# clear trace
VSP-9012:1(config)# trace spbm isis level 3
VSP-9012:1(config)# trace shutdown
VSP-9012:1(config)# show trace spbm isis
=====
                        SPBM ISIS Tracing Info
=====
Status      : Enabled
Level       : VERY_TERSE
Flag Info   :
VSP-9012:1(config)#trace level 125 3
VSP-9012:1(config)#trace level 119 3
VSP-9012:1(config)# save trace
VSP-9012:1(config)# trace grep error
VSP-9012:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable definitions

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

Table 35: Variable definitions

Variable	Value
<code>grep [WORD<0-128>]</code>	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
<code>level <0-192>[<0-4>]</code>	Starts the trace by specifying the module ID and level. <ul style="list-style-type: none"> • <code><0-192></code> specifies the module ID. • <code><0-4></code> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose.
<code>spbm isis level [<0-4>]</code>	Starts the trace by specifying the level. <ul style="list-style-type: none"> • <code><0-4></code> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is very verbose, 4 is verbose. <p>The default is 1, very terse.</p>
<code>shutdown</code>	Stops the trace operation.
<code>screen {disable enable}</code>	Enables the display of trace output to the screen.

Use the data in the following table to use the `save trace` command.

Table 36: Variable definitions

Variable	Value
<code>file WORD<1-99></code>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> • <code>a.b.c.d:<file></code> • <code>x:x:x:x:x:x:x <file></code> • <code>/intflash/<file></code> • <code>/extflash/<file></code> • <code>/usb/<file></code> • <code>/mnt/intflash/ <file></code> • <code>/mnt/extflash/ <file></code> <p><code>/mnt/intflash</code> is the internal flash of the second CP module (the one to which you are not connected).</p> <p><code>/mnt/extflash</code> is the external flash of the second CP module (the one to which you are not connected).</p>

CFM configuration using EDM

This section provides procedures to configure Connectivity Fault Management (CFM) using Enterprise Device Manager (EDM).

*** Note:**

When you enable CFM in an SPBM network, Avaya recommends that you enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

Autogenerated CFM

CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure an MD, MA, and MEP ID to create a MEP.

*** Note:**

- For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN. VSP 9000 only supports one MEP or MIP on the SPBM B-VLAN, either explicitly configured or autogenerated.

For autogenerated CFM configuration information for EDM see:

- [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 224.
- [Configuring autogenerated CFM on C-VLANs](#) on page 226.

Configuring autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This eliminates the need to explicitly configure an MD, MA, and MEP ID and to associate the MEP and MIP level to the SPBM B-VLAN.

To configure autogenerated CFM on C-VLANs, see [Configuring autogenerated CFM on C-VLANs](#) on page 226.

About this task

When you enable this feature, the device creates a global MD (named spbm) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

! Important:

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **Global** tab.
4. Click **enable** next to **SpbmAdminState**.
5. Specify a maintenance level in the **SPBMLevel** field.
6. Specify an MEP ID in the **SpbmMepId** field.
7. Click **Apply**.

CFM Global field descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable. Enabling SpbmAdminState creates one MIP level and one MEP on every B-VLAN at the specified SpbmLevel.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each B-VLAN type.
SpbmMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
CmacAdminState	Enables or disables autogenerated CFM for C-VLANs. The default is disable. Enabling CmacAdminState creates one MIP level and one MEP on every C-VLAN at the specified CmacLevel
CmacLevel	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4.

Table continues...

Name	Description
	Only configure global CFM at one MD level for each chassis for each C-VLAN type.
CmacMepld	Specifies the global CFM CMAC MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
Bmac	This is read-only only field. Specifies the B-MAC address of the node.
Cmac	This is read-only only field. Specifies the C-MAC address of the node.

Configuring autogenerated CFM on C-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every C-VLAN on the chassis.

To configure autogenerated CFM on SPBM B-VLANs, see [Configuring autogenerated CFM on SPBM B-VLANs](#) on page 224.

Important:

CFM supports only one MEP and one MIP per C-VLAN. The only method to configure CFM MEPs and MIPs on C-VLANs is to use the simplified commands above which autogenerated MEPs and MIP on C-VLANs. You cannot explicitly configure on some VLANs as this is possible with SPBM CFM.

About this task

When you enable this feature, you create a global MD (named cmac) for all the customer MAC (C-MAC) MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. The autogenerated CFM commands also create an MA for each C-VLAN, a MEP for each C-VLAN, and associate the MEP with the corresponding C-VLAN and a MIP with the C-VLAN.

All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The device automatically associates the MEPs with the C-VLANs configured. The MIP level maps to the global level. The device automatically associates the MIP level with the C-VLANs configured when you enable the feature.

The feature is disabled by default.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **Global** tab.
4. Click **enable** next to CmacAdminState.
5. Specify a C-MAC CFM maintenance level in the **CmacLevel** field.
6. Specify an MEP ID in the **CmacMepld** field.

7. Click **Apply**.

CFM Global field descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
SpbmMepId	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
CmacAdminState	Enables or disables autogenerated CFM for C-VLANs. The default is disable.
CmacLevel	Specifies the global C-MAC CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
CmacMepId	Specifies the global CFM MEP within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
Bmac	Specifies the B-MAC address of the node.
Cmac	Specifies the C-MAC address of the node.

Configuring explicit CFM

For SPBM B-VLANs, CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. Use the procedures in this section to configure MEPs explicitly.

If you want to create autogenerated CFM MEPs that eliminate the need to configure an MD, MA, and MEP ID, see the procedures in [Autogenerated CFM](#) on page 224. For C-VLANs, you can only use the autogenerated method.

Note:

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

For explicit configuration information for EDM see:

- [Configuring CFM MD](#) on page 228.
- [Configuring CFM MA](#) on page 228.
- [Configuring CFM MEP](#) on page 229.
- [Configuring CFM nodal MEP](#) on page 231.

Configuring CFM MD

Use this procedure to configure a Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

* Note:

If you use autogenerated CFM, you do not configure CFM MD because VSP 9000 configures a default MD, MA, MEPs, and MIPs.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **MD** tab.
4. Click **Insert**.
5. In the fields provided, specify an index value, name, and level for the MD.
6. Click **Insert**.

MD field descriptions

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

Name	Description
Index	Specifies a maintenance domain entry index.
Name	Specifies the MD name.
NumOfMa	Indicates the number of MAs that belong to this maintenance domain.
Level	Specifies the MD maintenance level. The default is 4.
NumOfMip	Indicates the number of MIPs that belong to this maintenance domain
Type	Indicates the type of domain.

Configuring CFM MA

Use this procedure to configure a CFM Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance

Endpoints (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

*** Note:**

If you use autogenerated CFM, you do not configure CFM MA because the switch configures a default MD, MA, MEPs, and MIPs.

Before you begin

- You must configure a CFM MD.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **MD** tab.
4. Highlight an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, click **Insert**.
6. In the fields provided, specify an index value and name for the MA.
7. Click **Insert**.

MA field descriptions

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

Name	Description
DomainIndex	Specifies the maintenance domain entry index.
AssociationIndex	Specifies a maintenance association entry index.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
NumOfMep	Indicates the number of MEPs that belong to this maintenance association.

Configuring CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

*** Note:**

If you use autogenerated CFM, you do not configure CFM MEPs because VSP 9000 configures a default MD, MA, MEPs, and MIPs.


Procedure

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

2. Click **CFM**.
3. Click the **MD** tab.
4. Highlight an existing MD, and then click **MaintenanceAssociation**.
5. In the **MA** tab, highlight an existing MA, and then click **MaintenanceEndpoint**.
6. Click **Insert**.
7. In the fields provided, specify the ID and the administrative state of the MEP.
8. Click **Insert**.

MEP field descriptions

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

Name	Description
DomainIndex	Specifies the MD index.
AssociationIndex	Specifies the MA index.
Id	Specifies the MEP ID.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
AdminState	Specifies the administrative state of the MEP. The default is disable.
MepType	<p>Specifies the MEP type:</p> <ul style="list-style-type: none"> • trunk • sg • endpt • vlan • port • endptClient • nodal • remotetrunk • remotesg • remoteendpt • remoteVlan • remotePort • remoteEndptClient <p> Note: VSP products only support Nodal Mep Type.</p>
ServiceDescription	Specifies the service to which this MEP is assigned.

Configuring CFM nodal MEP

Use this procedure to configure the CFM nodal Maintenance Endpoint (MEP). The Nodal MEP provides traceability and troubleshooting at the system level for a specific B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the specific B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and Maintenance Intermediate Point (MIP) functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a specific MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

* Note:

If you use autogenerated CFM, you do not configure CFM nodal MEPs because the switch configures a default MD, MA, MEPs, and MIPs.

Before you begin

- You must configure a CFM MD, MA, and MEP.

Procedure

- In the navigation tree, expand the following folders: **Configuration > VLAN**.
- Click **VLANs**.
- Click the **Advanced** tab.
- Select an SPBM VLAN.
- Click **Nodal**.
- In the **NodalMepList** field, specify the nodal MEPs to add to the VLAN.
- Click **Apply**.

Nodal MEP/MIP field descriptions

Use the data in the following table to use the **Nodal MEP/MIP** tab.

Name	Description
NodalMepList	Specifies the nodal MEPs to add to the VLAN, in the format <mdName.maName.mepId>, for example md10.ma20.30.
NumOfNodalMep	Indicates the number of nodal MEPs assigned to this VLAN.
NodalMipLevelList	Specifies a MIP level list.
NumOfNodalMipLevel	Indicates the number of nodal MIP levels assigned to this VLAN that allows MIP functionality to be enabled for each level for each VLAN.

Configuring Layer 2 ping

Use this procedure to configure a Layer 2 ping for C-VLANs or B-VLANs. This feature enables CFM to debug Layer 2. Layer 2 ping can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

*** Note:**

To use Layer 2 ping for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with **L2Ping**, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2Ping** option to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN or you can autogenerated CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerated CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. From the **L2Ping** tab, configure the Layer 2 ping properties.
4. To initiate a Layer 2 ping, highlight an entry and click the **Start** button.
5. To update a Layer 2 ping, click the **Refresh** button.
6. To stop the Layer 2 ping, click the **Stop** button.

L2Ping field descriptions

Use the data in the following table to use the **L2Ping** tab.

Name	Description
VlanId	Identifies the B-VLAN or the C-VLAN.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Ping transmission.
Messages	Specifies the number of L2Ping messages to be transmitted. The default is 1.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the L2Ping messages. • abort: The service aborted or is about to abort the L2Ping messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: The L2Ping Messages will be (or have been) sent. • false: The L2Ping Messages will not be sent. <p>The default is true.</p>
Priority	<p>Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame.</p> <p>The default is 7.</p>
TimeoutInt	Specifies the interval to wait for an L2Ping time-out. The default value is 3 seconds.
TestPattern	<p>Specifies the test pattern to use in the L2Ping PDU:</p> <ul style="list-style-type: none"> • allZero: Null signal without cyclic redundancy check. • allZeroCrc: Null signal with cyclic redundancy check with 32-bit polynomial. • pseudoRandomBitSequence: Pseudo-random-bit-sequence without cyclic redundancy check. • pseudoRandomBitSequenceCrc: Pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial.

Table continues...

Name	Description
	A cyclic redundancy check is a code that detects errors. The default value is allZero.
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The default is 0.
FrameSize	Specifies the frame size. If the frame size is specified then the data size is internally calculated and the calculated data size is included in the data TLV. The default is 0.
SourceMode	Specifies the source modes of the transmit loopback service: <ul style="list-style-type: none"> • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the B-MAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Result	Displays the Layer 2 Ping result.

Initiating a Layer 2 traceroute

Use this procedure for B-VLANs or C-VLANs to trigger a Layer 2 traceroute. This feature enables CFM to debug Layer 2. Layer 2 traceroute can also help you debug ARP problems by providing the ability to troubleshoot next hop when it modifies ARP records.

*** Note:**

To use Layer 2 traceroute for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify a MAC address with **L2Traceroute**, the MAC address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2Traceroute** option to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN, or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.

About this task

If you configure **IsTraceTree** to false, then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true, then EDM performs TraceTree on the multicast tree.

For more information on how to configure tracetable, see [Configuring Layer 2 tracetable](#) on page 252.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **L2Ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.
- For C-VLANs, you have to trigger an **L2Ping** to learn the C-MAC address.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 Traceroute/TraceTree** tab.
4. To start the traceroute, highlight an entry, and then click the **Start** button.
5. To update the traceroute, click the **Refresh** button.
6. To stop the traceroute, click the **Stop** button.

L2Traceroute field descriptions

Use the data in the following table to use the **L2Traceroute** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the B-VLAN or C-VLAN.
Priority	Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.

Table continues...

Name	Description
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Specifies whether the host name is (true) or is not (false) used for the L2Trace transmission.
Isid	Specifies the Service Instance Identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced. If you configure IsTraceTree to false then EDM performs Traceroute on the unicast path. If you configure IsTraceTree to true then EDM performs TraceTree on the multicast tree.
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the L2Trace messages. • abort: The service aborted or is about to abort the L2Trace messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: The L2Trace messages will be (or have been) sent. • false: The L2Trace messages will not be sent. <p>The default is true.</p>
Ttl	<p>Specifies the number of hops remaining to this L2Trace.</p> <p>This value is decremented by 1 by each bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the time-to-live (TTL) field in the L2Trace is defined by the originating MEP.</p> <p>The default value is 64.</p>
SourceMode	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • nodal • noVlanMac — Use this value with C-VLANs only. When you select this option, even if a VLAN MAC

Table continues...

Name	Description
	<p>address exists, the system uses the CFM C-MAC as the B-MAC-SA.</p> <ul style="list-style-type: none"> • <code>smltVirtual</code> — Use the <code>smltVirtual</code> option with B-VLANs only. <p>The default is <code>nodal</code>.</p>
SeqNumber	Specifies the transaction identifier/sequence number of the first linktrace message sent. The default is 0.
Flag	<p>L2Trace result flag indicating L2Trace status or error code:</p> <ul style="list-style-type: none"> • <code>none</code> (1): No error • <code>internalError</code> (2): L2Trace internal error • <code>invalidMac</code> (3): Invalid MAC address • <code>mepDisabled</code> (4): MEP must be enabled to perform L2Trace • <code>noL2TraceResponse</code> (5): No L2Trace response received • <code>l2TraceToOwnMepMac</code> (6): L2Trace to own MEP MAC is not sent • <code>l2TraceComplete</code> (7): L2Trace completed • <code>l2TraceLookupFailure</code> (8): Lookup failure for L2Trace • <code>l2TraceLeafNode</code> (9): On a leaf node in the I-SID tree • <code>l2TraceNotInTree</code> (10): Not in the I-SID tree • <code>l2TraceSmltNotPrimary</code> (11): Requested SMLT source from non-primary node

Viewing Layer 2 traceroute results

Use this procedure to view Layer 2 traceroute results. This feature enables CFM to debug Layer 2. It can also help you debug ARP problems by providing the ability to troubleshoot next hop ARP records.

About this task

You can display Layer 2 tracetable results to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID. For more information, see [Viewing Layer 2 tracetable results](#) on page 255.

Procedure

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

2. Click **L2Ping/L2Trace Route**.
3. Click the **L2Traceroute/TraceTree** tab.
4. Click the **Refresh** button to update the results.
5. To view the traceroute results, highlight an entry, and then click **Result**.

L2 Traceroute Result field descriptions

Use the data in the following table to use the **L2 Traceroute Result** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the B-VLAN or C-VLAN.
SeqNumber	Specifies the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which response of the L2Trace is going to be returned. The default is 0.
Hop	Specifies the number of hops away from L2Trace initiator.
ReceiveOrder	Specifies an index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Specifies a time-to-Live (TTL) field value for a returned L2Trace response.
SrcMac	Specifies the MAC address of the MP that responds to the L2Trace request for this L2TraceReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	Specifies the host name of the node that forwarded the L2Trace to the responding node.

Configuring Layer 2 IP ping

Use this procedure for B-VLANs or C-VLANs to configure Layer 2 IP ping.

Layer 2 IP ping allows a user to specify an IP address as the destination address. In this case, the IP address can be for a B-VLAN or C-VLAN.

*** Note:**

To use Layer 2 IP Ping for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between

the devices. When you specify an IP address with **L2 IP Ping**, the IP address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2 IP Ping** option to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.
- If you want to run a Layer 2 IP ping for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in *Configuring IP Routing on Avaya Virtual Services Platform 9000*, NN46250-505.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Ping** tab.
4. To add a new entry, click **Insert**, specify the destination IP address and optional parameters, and click **Insert**.
5. To start the Layer 2 IP ping, highlight an entry, and then click **Start**.
6. To update the Layer 2 IP ping, click the **Refresh** button.
7. To stop the Layer 2 IP ping, click **Stop**.

L2 IP Ping field descriptions

Use the data in the following table to use the **L2 IP Ping** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address. Only IPv4 is supported.
IpAddr	Specifies the destination IP address.

Table continues...

Name	Description
Vrflid	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Messages	Specifies the number of L2IpPing messages to be transmitted for each MAC/VLAN pair. Range is 1–200. The default is 1.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the L2IpPing messages. • abort: The service is aborted or about to abort the L2IpPing messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: L2IpPing Messages will be or have been sent. • false: L2IpPing Messages will not be sent. <p>The default is true.</p>
TimeoutInt	Specifies the interval to wait for an L2IpPing time-out with a range of 1–10 seconds with a default value of 3 seconds.
TestPattern	<p>Specifies the test pattern to use in the L2IPPing PDU:</p> <ul style="list-style-type: none"> • allZero — Null signal without cyclic redundancy check. • allZeroCrc — Null signal with cyclic redundancy check with 32-bit polynomial. • pseudoRandomBitSequence — Pseudo-random-bit-sequence without cyclic redundancy check. • pseudoRandomBitSequenceCrc — Pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors.</p> <p>The default value is allZero.</p>

Table continues...

Name	Description
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The range is 0–400. The default is 0.
PathsFound	Specifies the number of paths found to execute the command. The default is 0.

Viewing Layer 2 IP ping results

Use this procedure to view Layer 2 IP ping results.

*** Note:**

After you trigger Layer 2 IP ping, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 IP Ping** tab.
4. To view the Layer 2 IP ping results, highlight an entry, and then click **Result**.

L2 IP Ping Result field descriptions

Use the data in the following table to use the **L2 IP Ping Result** tab.

Name	Description
IpAddrType	Specifies the address type of the destination IP address.
IpAddr	Specifies the destination IP address.
SendOrder	Specifies the order that sessions were sent. It is an index to distinguish among multiple L2Ping sessions. This value is assigned sequentially from 1. It correlates to the number of paths found.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN ID found from the Layer 3 lookup and used for transmission.
DestMacAddress	Indicates the target MAC address transmitted.
PortNum	Specifies either the value '0', or the port number of the port used for the L2 IP ping.
DestHostName	Specifies the host name of the responding node.
Size	Specifies the number of bytes of data sent.

Table continues...

Name	Description
PktsTx	Specifies the number of packets transmitted for this VLAN/MAC.
PktsRx	Specifies the number of packets received for this VLAN/MAC.
PercentLossWhole	Specifies the percentage of packet loss for this VLAN/MAC.
PercentLossFract	Specifies the percentage of packet loss for this VLAN/MAC.
MinRoundTrip	Specifies the minimum time for round-trip for this VLAN/MAC.
MaxRoundTrip	Specifies the maximum time for round-trip for this VLAN/MAC.
RttAvgWhole	Specifies the average time for round-trip for this VLAN/MAC.
RttAvgFract	Specifies the fractional portion of average time for round-trip.
Flag	Specifies the result flag indicating status or error code: <ul style="list-style-type: none"> • 1 - No error • 2 - Internal error • 3 - Invalid IP • 4 - L2Trace completed • 5 - Lookup failure for IP (no VLAN/MAC entries)

Configuring Layer 2 IP traceroute

Use this procedure for C-VLANs or B-VLANs to configure Layer 2 IP traceroute.

*** Note:**

To use Layer 2 IP traceroute for C-MAC addresses to test connectivity between access points, you must configure autogenerated CFM and configure a traditional VLAN or a Layer 2 VSN between the devices. When you specify an IP address with **L2 IP Traceroute**, the IP address must be reachable on the VLAN specified on the command line.

Use SPBM CFM to test connectivity on the B-VLANs within an SPBM network. You can use the **L2 IP Traceroute** option to test reachability for all the B-MAC addresses in the SPBM network.

*** Note:**

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. Once you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **L2Ping**, **L2 IP Ping**, **L2 Traceroute**, and **L2 IP Traceroute** requests.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must either explicitly configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN or you can autogenerate CFM MEP and MIP for B-VLAN.
- If you want to do a Layer 2 ping for a C-VLAN, you must autogenerate CFM MEP and MIP for the C-VLAN. You cannot explicitly configure CFM MD, MA, and MEP for C-VLANs.
- If you want to run a Layer 2 IP traceroute for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in *Configuring IP Routing on Avaya Virtual Services Platform 9000*, NN46250-505.

About this task

If you configure **IsTraceTree** to false, then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true, then EDM performs TraceTree on the multicast tree.

For more information on how to configure tracetree, see [Configuring Layer 2 tracetree](#) on page 252.

Important:

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **L2Ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.
- For C-VLANs, you have to trigger an **L2Ping** to learn the C-MAC address.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**
3. Click the **L2 IP Traceroute** tab.
4. To add a new entry, click **Insert**, specify the destination IP address and, optionally, the TTL value, and then click **Insert**.
5. To start the Layer 2 IP traceroute, highlight an entry, and then click the **Start** button.
6. To update the L2 IP traceroute, click the **Refresh** button.
7. To stop the Layer 2 IP traceroute, click the **Stop** button.

L2 IP Traceroute field descriptions

Use the data in the following table to use the **L2 IP Traceroute** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address. Only IPv4 is supported.
IPAddr	Specifies the destination IP address.
Vrflid	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The default value is 64.
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit, the L2Trace messages. • abort: Specifies the service is aborted or about to abort the L2Trace messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time. The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the Trace Messages will be or have been sent. • false: the Trace Messages will not be sent <p>The default is true.</p>
PathsFound	Specifies the number of paths found to execute the L2Trace. The default is 0.

Viewing Layer 2 IP traceroute results

Use this procedure to view Layer 2 IP traceroute results.

*** Note:**

After you trigger Layer 2 IP traceroute, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.

3. Click the **L2 IP Traceroute** tab.
4. To view the Layer 2 IP traceroute results, highlight an entry, and then click **Result**.

L2 IP Traceroute Result field descriptions

Use the data in the following table to use the **L2 IP Traceoute Result** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address.
IpAddr	Specifies the destination IP address.
SendOrder	Denotes the order that sessions are sent. It is an index to distinguish among multiple L2Trace sessions. It correlates to the number of paths found. This value is assigned sequentially from 1.
Hop	Specifies the number of L2 hops away from L2Trace initiator.
ReceiveOrder	Specifies the order that sessions are sent. It is an index to distinguish among multiple L2Trace responses with the same Send Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Specifies the time-to-live (TTL) field value for a returned L2Trace response.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN found from Layer 3 lookup and used for transmission.
DestMacAddress	Indicates the target MAC address transmitted.
PortNum	Specifies either the value '0', or the port number of the port used for the l2trace.
SeqNumber	Specifies the transaction identifier/sequence number used in linktrace message packet. The default is 0.
SrcMac	Specifies the MAC address of the MP that responded to L2Trace request for this L2traceReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	Specifies the host name of the node that forwarded the L2Trace to the responding node.
Flag	Indicates the L2Trace result flag status or error code: <ul style="list-style-type: none"> • none (1): No error • internalError (2): L2Trace internal error

Table continues...

Name	Description
	<ul style="list-style-type: none"> • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace

Triggering a loopback test

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN or C-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **LBM** tab.
4. Configure the loopback test properties as required.
5. Click **Apply**.
6. To trigger the loopback test, double-click in the **Status** field, select **transmit**.
7. Click **Apply**.
8. To update the loopback test, click the **Refresh** button.

LBM field descriptions

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

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the Maintenance Endpoint index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
DestMacAddress	Specifies the remote MAC address to reach the MEP/MIP.
Messages	Specifies the number of loopback messages to be transmitted. The default is 1.
VlanPriority	Specifies the priority. The default is 7.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: <ul style="list-style-type: none"> • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.
Status	Indicates the status of the transmit loopback service: <ul style="list-style-type: none"> • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the Loopback messages. • abort: The service is aborted or about to abort the Loopback messages. The default is ready.
Result	Displays the LBM result.
TimeoutInt	Specifies the timeout interval in seconds. The default value is 3 seconds.
InterFrameInt	Specifies the interval between LBM frames with a range of (0..1000) msec and a default value of 500 msec. The value of 0 msec indicates to send the frames as fast as possible. The default is 500.
TestPattern	Specifies the testfill pattern: <ul style="list-style-type: none"> • allZero — null signal without cyclic redundancy check • allZeroCrc — null signal with cyclic redundancy check with 32-bit polynomial

Table continues...

Name	Description
	<ul style="list-style-type: none"> • pseudoRandomBitSequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudoRandomBitSequenceCrc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial <p>A cyclic redundancy check is a code that detects errors. The default value is allZero.</p>
DataSize	Specifies the data type-length-value (TLV) size. The default is 0.
FrameSize	Specifies the frame-size. The default is 0.
SourceMode	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the B-MAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. <p>The default is nodal.</p>

Triggering linktrace

Use the following procedure to trigger a linktrace. The link trace message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Before you begin

- You must configure and enable CFM.
- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN or C-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **LTM** tab.
4. Configure the linktrace test properties as required.
5. Click **Apply**.

6. To trigger the linktrace test, double-click in the Status field, select **transmit**, and then click **Apply**.

OR

Highlight an entry, and then click **Start**.

7. To update the linktrace, click the **Refresh** button.
8. To stop the linktrace, click **Stop**.
9. To view the results of the linktrace, click **Result**.

LTM field descriptions

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

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the MEP index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
VlanPriority	Specifies the VLAN priority, a 3-bit value to be used in the VLAN tag, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the remote MAC address to reach the MEP.
Ttl	Indicates the number of hops remaining to this LTM. This value is decremented by 1 by each bridge that handles the LTM. The decremented value is returned in the LTR. If the value is 0 on output, the LTM is not transmitted to the next hop. The value of the TTL field in the LTM is specified at the originating MEP. The default value is 64.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: <ul style="list-style-type: none"> • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.
Status	Indicates the status of the transmit loopback service: <ul style="list-style-type: none"> • ready: The service is available.

Table continues...

Name	Description
	<ul style="list-style-type: none"> • transmit: The service is transmitting, or about to transmit, the LTM messages. • abort: The service is aborted, or about to abort, the LTM message. <p>The default is ready.</p>
Flag	<p>Displays the LTM result flag indicating LTM status or error code. Each value represents a status or error case:</p> <ul style="list-style-type: none"> • 1 - No error • 2 - LTM internal error • 3 - Unknown Remote Maintenance Endpoint • 4 - Invalid Remote Maintenance Endpoint MAC Address • 5 - Unset Remote Maintenance Endpoint MAC address • 6 - MEP must be enabled to perform LTM • 7 - No LTR response received • 8 - Linktrace to own MEP MAC is not sent • 9 - Endpoint must be enabled in order to perform LTM • 10 - Pbt-trunk must be enabled to perform LTM • 11 - LTM completed • 12 - LTM leaf node
SourceMode	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • nodal • noVlanMac — Use this value with C-VLAN only. When you select this option, even if a VLAN MAC address exists, the system uses the CFM C-MAC as the BMAC-SA. • smltVirtual — Use the smltVirtual option with B-VLANs only. <p>The default is nodal.</p>

Viewing linktrace results

Use this procedure to view linktrace results.

*** Note:**

After you trigger linktrace, you must click the **Refresh** button to update the results.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **CFM**.
3. Click the **LTM** tab.
4. Highlight an entry, and then click **Result**.

Link Trace Replies field descriptions

Use the data in the following table to use the **Link Trace Result** tab.

Name	Description
DomainIndex	Indicates the Maintenance Domain Index.
AssociationIndex	Indicates the Maintenance Association Index.
MepId	Indicates the Maintenance EndPoint ID.
SeqNumber	Indicates the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which LTM response is going to be returned. The default is 0.
Hop	Indicates the number of hops away from the LTM initiator.
ReceiveOrder	Indicates the index value used to distinguish among multiple LTRs with the same LTR Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the LTRs.
Ttl	Indicates the TTL field value for a returned LTR.
DomainName	Indicates the Maintenance Domain Name.
AssociationName	Indicates the Maintenance Association Name.
Forwarded	Indicates if a LTM was forwarded by the responding MP, as returned in the FwdYes flag of the flags field.
TerminalMep	Displays a boolean value stating whether the forwarded LTM reached a MEP enclosing its MA, as returned in the Terminal MEP flag of the Flags field.
LastEgressIdentifier	Displays an octet field holding the Last Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Last Egress Identifier identifies the MEP Linktrace Indicator that originated, or the Linktrace Responder that forwarded, the LTM to

Table continues...

Name	Description
	which this LTR is the response. This is the same value as the Egress Identifier TLV of that LTM.
NextEgressIdentifier	Displays an octet field holding the Next Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Next Egress Identifier Identifies the Linktrace Responder that transmitted this LTR, and can forward the LTM to the next hop. This is the same value as the Egress Identifier TLV of the forwarded LTM, if any. If the FwdYes bit of the Flags field is false, the contents of this field are undefined, and the field is ignored by the receiver.
RelayAction	Indicates the value returned in the RelayAction field.
SrcMac	Displays the MAC address of the MP that responded to the LTM request for this LTR.
IngressAction	Displays the value returned in the IngressAction Field of the LTM. The value ingNoTlv indicates that no Reply Ingress TLV was returned in the LTM.
IngressMac	Displays the MAC address returned in the ingress MAC address field. If the rcCfmLtrReplyIngress object contains the value ingNoTlv(5), then the contents of this field are meaningless.
EgressAction	Displays the value returned in the Egress Action Field of the LTM. The value egrNoTlv(5) indicates that no Reply Egress TLV was returned in the LTM.
EgressMac	Displays the MAC address returned in the egress MAC address field. If the rcCfmLtrReplyEgress object contains the value egrNoTlv(5), then the contents of this field are meaningless.

Configuring Layer 2 tracetree

Use this procedure to configure a Layer 2 Tracetree. This feature enables CFM to debug Layer 2. Layer 2 Tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.

Note:

VSP 9000 only supports this command on SPBM B-VLANs only, not C-VLANs.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.

- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. From the **L2 Traceroute/TraceTree** tab, configure the Layer 2 tracetree properties.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Click **Apply**.
6. Click the **Refresh** button to update the results.

L2Tracetree field descriptions

Use the data in the following table to use the **L2Tracetree** tab.

Name	Description
VlanId	Identifies the Backbone VLAN.
Priority	Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Tracetree transmission.
Isid	Specifies the service instance identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced. If you configure IsTraceTree to false then EDM performs Traceroute on the unicast path. If you configure IsTraceTree to true then EDM performs TraceTree on the multicast tree.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Tracetree messages. • abort: the service aborted or is about to abort the L2Tracetree messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p>

Table continues...

Name	Description
	The default is ready.
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the L2Tracetree Messages will be (or have been) sent. • false: the L2Tracetree Messages will not be sent <p>The default is true.</p>
Ttl	<p>Specifies the Time-to-Live value. Indicates the number of hops remaining to this L2Tracetree. The tracetree is decremented by one by each bridge that handles the Layer 2 tracetree and the decremented value is returned to the tracetree. If the output is 0, then the L2Tracetree is not transmitted to the next hop. The value of the TTL field in the L2Tracetree is transmitted by the originating MEP is controlled by a managed object. The default is 64.</p>
SourceMode	<p>Specifies the source modes of the transmit loopback service:</p> <ul style="list-style-type: none"> • nodal • noVlanMac — Use the noVlanMac option with C-VLANs only. • smltVirtual — Use the smltVirtual option with B-VLANs only. <p>The default is nodal.</p>
SeqNumber	<p>The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.</p>
Flag	<p>Specifies the L2Tracetree result flag, which indicates the L2Tracetree status or error code. Each sum represents a status or error:</p> <ul style="list-style-type: none"> • 1 — No error • 2 — L2Tracetree internal error • 3 — Invalid MAC address • 4 — MEP must be enabled to perform L2Tracetree • 5 — No L2Tracetree response received • 6 — L2Tracetree to own MEP MAC is not sent • 7 — L2Tracetree completed • 8 — Lookup failure for L2Tracetree • 9 — On a leaf node in the I-SID tree • 10 — Not in the I-SID tree

Table continues...

Name	Description
	<ul style="list-style-type: none"> 11 — Requested SMLT source from non-primary node

Viewing Layer 2 tracetable results

Use this procedure to view Layer 2 Tracetable results. The Layer 2 Tracetable command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics**.
2. Click **L2Ping/L2Trace Route**.
3. Click the **L2 Traceroute/TraceTree** tab.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetable on the multicast tree.
5. Click **Apply**.
6. Click the **Refresh** button to update the results.
7. To view the tracetable results, highlight an entry, and then click **Result**.

L2 Tracetable Result field descriptions

Use the data in the following table to use the **L2 Tracetable Result** tab.

Name	Description
VlanId	A value that uniquely identifies the Backbone VLAN (B-VLAN).
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which response of the L2Tracetable is going to be returned. The default is 0.
Hop	The number of hops away from L2Tracetable initiator.
ReceiveOrder	An index to distinguish among multiple L2Tracetable responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Time-to-Live (TTL) field value for a returned L2Tracetable response.

Table continues...

Name	Description
SrcMac	MAC address of the MP that responds to the L2Tracetree request for this L2tractreeReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Tracetree to the responding node.
LastHostName	The host name of the node that forwarded the L2Tracetree to the responding node.

Configuring Layer 2 trace multicast route on a VLAN

Use this procedure to configure the Layer 2 tracemroute on the VLAN (Layer 2). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID, and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

* Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context. See the following procedure to perform a Layer 3 tracemroute on a VRF, [Configuring Layer 2 tracemroute on a VRF](#) on page 258.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics > L2Ping/L2Trace Route**.
2. Click the **L2MCAST Traceroute** tab.
3. Click **Insert** to insert the L2 MCAST Traceroute.
4. Type the **SrcIpAddr**.
5. Type the **GroupIpAddr**.
6. Enter the **ServiceType**. If you want to perform a Layer 2 tracemroute on a VLAN, select **vlan**. If you want to perform a Layer 2 tracemroute on a Layer 3 GRT, select **vrfid**.

* Note:

If you want to perform a Layer 2 tracemroute on a Layer 2 or a Layer 3 VRF, review the following procedure [Configuring Layer 2 tracemroute on a VRF](#) on page 258.

7. In the **ServiceId**, enter the VLAN ID.
8. Enter the **Priority**.
9. Enter the **Ttl** value.

10. Click **Insert**.
11. Click **Apply** to save your changes.
12. To start the Layer 2 tracemoute, set the Status to transmit and click the **Start** button.
13. Update the Layer 2 tracemroute by clicking the **Refresh** button.
14. To stop the Layer 2 tracemroute, click the **Stop** button.
15. To see the result, click the **Result** button.

L2 MCAST Traceroute field descriptions

Use the data in the following table to use the **L2MCAST Traceroute** tab.

Name	Description
SrcIpAddrType	Specifies the source IP address type as IPv4.
SrcIpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
VRFName	Specifies the VRF name.
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit the trace messages. • abort: Specifies the services is aborted or about to abort the trace messages. <p>The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p>
ResultOK	<p>Specifies the result of the operation:</p> <ul style="list-style-type: none"> • true: The trace messages will be or have been sent. • false: The trace messages will not be sent.

Table continues...

Name	Description
Flag	<p>Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case.</p> <ul style="list-style-type: none"> • 1 — No error • 2 — Internal Error • 3 — Mep must be enabled to perform the trace • 4 — No response received • 5 — Trace completed • 6 — On a leaf node in the I-SID tree • 7 — No data I-SID was found for S, G

Configuring Layer 2 tracemroute on a VRF

Use this procedure to configure the Layer 2 tracemroute on the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

* Note:

If you want to run a Layer 2 tracemroute on a VRF, make sure you are in the proper VRF context.

See the following procedure to perform a Layer 3 tracemroute on a VLAN [Configuring Layer 2 tracemroute on a VLAN](#) on page 256.

Before you begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > VRF Context View > Set VRF Context View**
2. Select a VRF and click the **Launch VRF Context View** tab.
3. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics > L2Ping/L2Trace Route**.
4. Click the **L2MCAST Traceroute** tab.
5. Click **Insert** to insert the L2 MCAST traceroute.
6. Type the **SrclpAddr**.
7. Type the **GroupIpAddr**.
8. Enter the **ServiceType**. If you want to perform a Layer 2 tracemroute on a Layer 2 VRF, select **vlan**. If you want to perform a Layer 2 tracemroute on a Layer 3 VRF, select **vrfid**.

9. In the **ServiceId**, enter the VLAN ID.
10. Enter the **Priority**.
11. Enter the **Ttl** value.
12. Click **Insert**.
13. Click **Apply** to save your changes.
14. To start the Layer 2 tracemoute, set the Status to transmit and click the **Start** button.
15. Update the Layer 2 tracemroute by clicking the **Refresh** button.
16. To stop the Layer 2 tracemroute, click the **Stop** button.
17. To see the result, click the **Result** button.

L2 MCAST Traceroute field descriptions

Use the data in the following table to use the **L2MCAST Traceroute** tab.

Name	Description
SrclpAddrType	Specifies the source IP address type as IPv4.
SrclpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
VRFName	Specifies the VRF name.
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: Specifies the service is available. • transmit: Specifies the service is transmitting, or about to transmit the trace messages. • abort: Specifies the services is aborted or about to abort the trace messages. <p>The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p>
ResultOK	<p>Specifies the result of the operation:</p> <ul style="list-style-type: none"> • true: The trace messages will be or have been sent.

Table continues...

Name	Description
	<ul style="list-style-type: none"> • false: The trace messages will not be sent.
Flag	<p>Specifies the result flag indicating that the L2 trace status or error code. Each value represents a status or error case.</p> <ul style="list-style-type: none"> • 1 — No error • 2 — Internal Error • 3 — Mep must be enabled to perform the trace • 4 — No response received • 5 — Trace completed • 6 — On a leaf node in the I-SID tree • 7 — No data I-SID was found for S, G

Viewing Layer 2 trace multicast route results

Use this procedure to view Layer 2 tracemroute results.

Procedure

1. From the navigation tree, expand the following folders: **Configuration > Edit > Diagnostics > L2Ping/L2Trace Route**
2. Click the **L2 MCAST Traceroute** tab.
3. To view the CFMI2 trace multicast route results, highlight an entry and click the **Result** button.

L2tracemroute Result field descriptions

Use the data in the following table to use the **L2tracemroute Result** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the C-VLAN.
SeqNumber	Specifies the transaction identifier/sequence number returned by a previous transmit linktrace message command. Indicates which I2 tracemroute response is going to be returned.
Hop	Specifies the number of hops away from the I2 tracemroute initiator.
ReceiveOrder	Specifies an index to distinguish among multiple I2 tracemroute responses with the same transaction identifier field value. This value is assigned sequentially from 1, in the order that the linktrace initiator received the responses.
Ttl	Specifies the TTL value for a returned I2 tracemroute response.

Table continues...

Name	Description
SrcMac	Specifies the MAC address of the MP that responds to the I2 tracemroute request for this I2 tracemrouteReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the I2 tracemroute to the responding node.
LastHostName	Specifies the host name of the node that forwarded the I2 tracemroute to the responding node.

Chapter 11: Upper layer troubleshooting

This section describes troubleshooting for Layer 4 to 7 applications.

Troubleshooting SNMP

About this task

Troubleshoot Simple Network Management Protocol (SNMP) if the network management station (NMS) does not receive traps.

Verify the management configurations for the management station. Also verify the management station setup. If the management station can reach a device but not receive traps, verify the trap configurations (that is, the trap destination address and the traps to be sent).

Procedure

1. From the NMS, ping the IP address for the switch. If you can ping successfully, the IP address is valid and you may have a problem with the SNMP setup.
If you cannot ping the switch, you have a problem with either the path or the IP address.
2. Telnet to the switch.
If you can Telnet, the switch IP address is correct.
3. If Telnet does not work, connect to the console port using a serial line connection and ensure that the IP address configuration is correct.
4. If the management station is on a separate subnet, make sure that the gateway address and subnet mask are correct.
5. Using a management application, perform an SNMP Get request and an SNMP Set request (that is, try to poll the device or change a configuration using management software).
6. If you cannot reach the device using SNMP, access the console port, and then ensure that the SNMP community strings and traps are correct.
7. Use sniffer traces to verify that the switch receives the poll.
8. Use sniffer traces to verify that the NMS receives the response.
9. Verify that the data in the response is the data that was requested.

Troubleshooting DHCP

About this task

Perform this procedure to troubleshoot the following Dynamic Host Configuration Protocol (DHCP) scenarios:

- The client cannot obtain a DHCP address when in the same subnet.
- The client cannot obtain a DHCP address when in a different subnet.

When the DHCP server and client are on the different subnets or VLANs, you must configure the device as a DHCP relay agent. The device must forward DHCP requests to the DHCP server. You must perform extra troubleshooting steps to troubleshoot the DHCP relay agent.

Procedure

1. Check the physical connectivity between the DHCP client and server.
2. Verify network connectivity by configuring a static IP address on a client workstation.
If the workstation still cannot reach the network, the problem is not DHCP. Start troubleshooting network connectivity.
3. Attempt to obtain an IP address from the DHCP server by manually forcing the client to send a DHCP request.
If the client obtains an IP address after the PC startup is complete, the issue is not the DHCP server.
4. Obtain an IP address on the same subnet or VLAN as the DHCP server.
If the issue persists, the problem may be with the DHCP server. If DHCP is working on the same subnet or VLAN as the DHCP server, the DHCP issue can be with the DHCP relay agent.
5. Confirm the DHCP relay agent configuration is correct.
6. Obtain sniffer traces where the traffic ingresses and egresses the switch and also on the client side of the network.
7. Check the logs on the switch for errors such as size exceeded or incorrect packet format.

Troubleshooting DHCP Relay

Before you begin

- Configure the server to reply to the client subnet. Check the server configuration file to verify the configuration.
- Configure a route on the server for the client subnet to create a path on which to send replies.

About this task

Perform this procedure to troubleshoot the DHCP relay agent.

Procedure

1. Verify that the interfaces that link the client and server are up, and that the ports are in the forwarding state.

- a. To verify client availability, you can configure a temporary static IP address on the client, and then use the `ping` command.

```
ping WORD<0-256>
```

- b. To verify the port is in the forwarding state, use the following command for the slot and port number:

```
show spanning-tree [rstp|mstp] port role [{slot/port[-slot/port]}
[,...]]
```

If STP detects loops in the configuration, it blocks ports to avoid flooding in the network. In this situation, the port is not in the forwarding state.

2. Ensure that DHCP is enabled on the client interface and that a valid forwarding path exists and is enabled. Ensure the server is reachable.
3. View the statistics counters for the relay.
4. If request or reply counters do not increase, use a sniffer tool to ensure that the client sends the packets, and that the interface module receives the packets.

You can configure mirroring for the ingress port to verify if the packets reach the module.

- a. If the client sends the packets, check that the packets reach the CPP and search the trace results for the ingress port:

```
trace level 9 3
trace grep WORD<0-128>
```

- b. If the packets reach the CPP, check that they reach the DHCP protocol; check for errors or packet drop messages:

```
trace level 170 3
trace grep WORD<0-128>
```

5. If Option 82 is enabled, check the statistic counters for dropped packets, and perform a trace for the DHCP protocol:

```
trace level 170 3
```

Example

To verify client availability, you can configure a temporary static IP address on the client, and then use the `ping` command. To verify the port is in the forwarding state, use the `show spanning-tree` command. If the client sends the packets, check that the packets reach the CPP and search the trace results for the ingress port.

```
VSP-9012:1>enable
VSP-9012:1#ping 192.2.0.2
VSP-9012:#show spanning-tree mstp port role
=====
CIST Port Roles and States
```



```

=====
Port-Index  Port-Role  Port-State  PortSTPStatus  PortOperStatus
-----
4/1         Disabled  Forwarding  Disabled       Disabled
4/2         Disabled  Forwarding  Disabled       Disabled
4/3         Disabled  Discarding  Enabled        Disabled
4/4         Disabled  Discarding  Enabled        Disabled
4/5         Disabled  Forwarding  Disabled       Disabled
4/6         Disabled  Forwarding  Disabled       Disabled
4/7         Disabled  Forwarding  Disabled       Disabled
4/8         Disabled  Forwarding  Disabled       Disabled
4/9         Disabled  Discarding  Enabled        Disabled
4/10        Disabled  Discarding  Enabled        Disabled
4/11        Disabled  Discarding  Enabled        Disabled
4/12        Designated Forwarding  Enabled        Enabled
4/13        Disabled  Forwarding  Disabled       Disabled
4/14        Disabled  Forwarding  Disabled       Disabled
4/15        Disabled  Discarding  Enabled        Disabled
4/16        Disabled  Discarding  Enabled        Disabled
4/17        Disabled  Discarding  Enabled        Disabled

--More-- (q = quit)

VSP-9012:1#trace level 9 3
VSP-9012:1#trace grep 00-1A-4B-8A-FB-6B

```

Variable definitions

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

Table 37: Variable definitions

Variable	Value
<code>WORD<0-256></code>	Specifies the IP address.


Use the data in the following table to use the `show spanning-tree` command.

Table 38: Variable definitions

Variable	Value
<code>{mstp rstp}</code>	Specifies the spanning tree protocol.
<code>port role{slot/port[-slot/port][,...]}</code>	Displays the port role information. Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).

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

Table 39: Variable definitions

Variable	Value
level [<0–192>][<1–4>]	<p>Starts the trace by specifying the module ID and level. <0-192> specifies the module ID.</p> <p><1–4> specifies the trace level:</p> <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable enable}	<p>Enables or disables the display of trace output to the screen.</p> <p> Important:</p> <p>Avaya recommends you avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.</p>

Troubleshooting client connection to the DHCP server

About this task

Perform this procedure if the client cannot reach the DHCP server.

Procedure

1. Check that the DHCP relay agent in the network switch is correctly configured.
2. Check that the DHCP server configuration is correct.
3. Check for routing issues.

The routing in the network may not be configured so that the DHCP request and reply packets are propagated. You can use ping and traceroute.

4. Check that the DHCP pools are correctly configured.
5. If the client cannot reach the server because the link is down, enable auto-negotiation on the link.

Troubleshooting IPv6 DHCP Relay

The following sections provide troubleshooting information for IPv6 DHCP Relay.

IPv6 DHCP Relay switch side troubleshooting

With DHCP Relay, the switch only participates in forwarding the requests and replies to and from the client and the DHCP server. The switch always acts as the relay agent, on which you configure the forward path to the server.

To troubleshoot DHCP Relay issues on the switch, use the following procedure.

Procedure

1. Verify that the DHCP server is reachable using ping. If ping is working and the DHCP server is reachable, DHCP should work.
2. Verify that the relay agents and the forward path configured are reachable. Ping the server and the gateway to the server.
3. Check that the relay agent configurations are correct. Also verify that DHCP is enabled on the switch:

```
show ipv6 dhcp-relay interface {gigabitEthernet {slot/port[-slot/
port][, ...]}|vlan <1-4084>
```

4. Verify that IPv6 forwarding is enabled globally:

```
show ipv6 global
```

5. Verify that the IPv6 based VLAN where the DHCP relay agent is configured is enabled:

```
show ipv6 interface vlan <1-4084>
```

6. In a scenario with VRRP and SMLT, Avaya recommends that you have the VRRP IP configured as the DHCP relay agent.

7. When using the VRRP VRID as the relay agent, make sure the VRRP configurations are proper.

8. To verify that relay forward and relay receive are working, enable trace for DHCP with IPv6, and grep trace for relay:

```
trace level 66 3
```

```
trace grep relay
```

```
trace screen enable
```

9. Display the count of DHCP Relay requests and replies to verify the system received requests and replies:

```
show ipv6 dhcp-relay counters
```

IPv6 DHCP Relay server side troubleshooting

Use the following procedure to troubleshoot IPv6 DHCP Relay on the server side.

Procedure

1. Enable the services on the server side, and then create an IP pool.
The IP pool must contain the range of addresses that you want to assign to the clients.
Configure the IP pool with the same network subnet as that of the relay agent.
2. When the configuration is complete, initiate a DHCP request from a client.
3. Check the log file available on the server to verify the reason for packet drop.
4. Capture the packets on the server side using Ethereal.
5. From the server side, use ping to verify that the relay agent address is reachable.
Ensure that a route to the relay is configured.
6. For more configuration aspects, see the MS webpage for troubleshooting and configuration issues.

* Note:

You can receive some log messages that indicate the system cannot forward packets. However, certain situations are not DHCP failures.

Example 1: if you receive the message `0x00108796 (relayMsgSend): cannot find route entry for destination` on the console, you must ping the server. If the server is not reachable, the system cannot forward the packet. This is not a DHCP issue.

Example 2: if you receive the message `0x00108705` this indicates a problem at the transmission level. Check the server reachability and ensure that MAC learning is correct before you pursue DHCP issues.

IPv6 DHCP Relay client side troubleshooting

You can collect a client console dump, which can be used to analyze why the received packet cannot be processed and the allocated address cannot be used by the client.

In addition, restarting the client can also fix the issue in some cases.

Make sure the client supports IPv6 requests.

Connect the server directly to the client. If the IP is assigned, then the problem is with the relay.

Enabling trace messages for IPv6 DHCP Relay

Use this procedure to enable trace for IPv6 DHCP Relay and enable IPv6 forwarding trace.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To troubleshoot IPv6 DHCP Relay, you can enable rcip6 trace messages using the following command:

```
trace level 66 3
```

3. You can also enable IPv6 forwarding trace using the following command:

```
trace ipv6 forwarding enable <all|debug|error|info|pkt|warn>
```

Example

Enable rcip6 trace messages and enable IPv6 forwarding trace:

```
Switch:1>enable
Switch:1#trace level 66 3
Switch:1#trace ipv6 forwarding
```

Troubleshooting IPv6 VRRP

The following sections describe troubleshooting information for IPv6 Virtual Router Redundancy Protocol (VRRP).

VRRP transitions

When a VRRP transition takes place with the backup taking over as the master, look for the following message in the syslog on the new master, as well as the old master. This message provides information to allow you to determine the cause of the transition.

```
IPv6 Vrrp State Transition Trap(Port/Vlan=200, Type=masterToInitialize,
Cause=shutdownReceived, VrId=20,VrIpAddr=fe80:0:0:0:0:0:0:200,
Addr=fe80:0:0:0:224:7fff:fe9d:1a03)
```

In this message, see the Type and Cause fields.

Note:

Although all of the possible causes and types are listed below, not all of the listed causes and types appear in the trap/log message.

The following table describes the VRRP transition types.

Table 40: Transition type

Type value	Type definition
1	None
2	Master to backup
3	Backup to master
4	Initialize to master
5	Master to initialize
6	Initialize to backup
7	Backup to initialize
8	Backup to backup master
9	Backup master to backup

The following table describes the VRRP transition causes.

Table 41: Transition cause

Cause value	Cause definition
1	None
2	Higher priority advertisement received
3	Shutdown received
4	VRRP address and physical address match
5	Master down interval
6	Preemption
7	Critical IP goes down
8	User disabling VRRP
9	VRRP status synced from primary
10	IPv6 interface on which VRRP is configured goes down
11	Lower priority advertisement received
12	Advertisement received from higher interface IP address with equal priority
13	Advertisement received from lower interface IP address with equal priority
14	User enabled VRRP
15	Transition because of any other cause

Enabling trace messages for IPv6 VRRP troubleshooting

Use this procedure to enable trace messages for IPv6 VRRP.

When VRRP is enabled on two routing switches, the master-backup relationship forms with one router taking the responsibility of routing. If the master-backup relationship is not formed between the VRRP virtual routers, look for the following trace messages to ensure that the master is sending the advertisements correctly and the backup is processing them.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. To troubleshoot IPv6 VRRP, you can enable RCIP6 trace messages with the command:

```
trace level 66 3
```

3. And to provide additional trace information, you can also enable the following traces:

```
trace ipv6 nd enable
```

```
trace ipv6 base enable all
```

```
trace ipv6 forwarding enable all
```

```
trace ipv6 rtm enable all
```

```
trace ipv6 transport enable all
```

4. When VRRP is enabled on two routing switches, the master-backup relationship forms with one router taking the responsibility of routing. If the master-backup relationship is not formed between the VRRP virtual routers, look for the following trace messages to ensure that the master is sending the advertisements correctly and the backup is processing them. On the master router, look for the following RCIP6 trace messages.

- tMainTask RCIP6: rcip6_vrrp.c: 5118: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Am Master for Vrid 200 on IfIndex 2053 Timer 1

If VRRP is enabled on the interface, this timer kicks off every second and shows the state for the VRID.

- [11/18/09 15:08:20:383] tMainTask RCIP6: rcip6_vrrp.c: 5924: ipv6VrrpSendAdvertisement: for Vrid 200 on IfIndex 2053

```
[11/18/09 15:08:20:583] tMainTask RCIP6: rcip6_vrrp.c: 5175: VRF
name: GlobalRouter (VRF id 0): ipv6VrrpTic:
ipv6VrrpSendAdvertisement
```

The preceding trace messages show that the VRRP master is sending the advertisements correctly at the end of advertisement interval for a VRID.

5. On the backup router, look for the following RCIP6 trace messages.

- tMainTask RCIP6: rcip6_vrrp.c: 5236: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Am Backup for VrId 200 on IfIndex 2052 Timer 1

- tMainTask RCIP6: rcip6_vrrp.c: 4854: ipv6VrrpIn: Vrid 200 on IfIndex 2052

- tMainTask RCIP6: rcip6_vrrp.c: 5545: VRF name: GlobalRouter (VRF id 0): rcIpVrrpProcessAdvt: Am backup for Vrid 200 on IfIndex 2052

The preceding trace messages show that the backup router is receiving the advertisements sent by the master and correctly processing them.

Risks associated with enabling trace messages

When traces are enabled on VRRP master, VrrpTic messages are logged for every second and any other configured traces keep displaying, so there is no guarantee that the backup will receive the advertisement from the master within 3 seconds, so it can transit to master also. There is also the risk of toggling of VRRP states (from backup to master and back again).

Enable the limited traces based on whichever is required.

VRRP with higher priority running as backup

The VRRP router with the higher priority can display as the backup for the following reasons

- Hold-down timer is running.
- The configured Critical IP is not reachable or does not exist.

If the critical-IP is configured for VRRP master, and the critical interface goes down or is deleted, the master transitions to the backup state. In this case, the log shows the transition cause as 1 like many other cases.

If the holddown timer is configured for VRRP master, the holddown timer delays the preemption, giving the device, which is becoming the master enough time to construct routing tables.

Procedure

1. To determine that the issue is with the critical interface, look for the following trace message.

```
tMainTask RCIP6: rcip6_vrrp.c: 5152: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Becoming backup for Vrid 200 on IfIndex 2052 because of invalid critical IP
```

2. If the holddown Timer is configured for VRRP master, the holddown timer delays the preemption, giving the device, which is becoming the master enough time to construct routing tables.

```
tMainTask RCIP6: rcip6_vrrp.c: Enter in HoldDown processing,Vrid 200 LastRecvd 0 MasterDown 3, Holddown time remaining 970, Holddownstate 2
```

Troubleshooting RSMLT

The following sections provide information for troubleshooting IPv4 and IPv6 Routed Split Multi-Link Trunking (RSMLT).

RSMLT configuration considerations

When troubleshooting IPv6 RSMLT, note the following configuration considerations:

- You must configure interswitch trunking (IST) peers with the same IPv6 subnets on the Split MultiLink Trunking (SMLT) VLANs (same as for IPv4).
- Make sure that the IST MultiLink Trunking (MLT) on the RSMLT peers contains the same set of links (this is very difficult to catch through regular troubleshooting).
- Running both IPv6 RSMLT and IPv6 VRRP on the same VLAN is not supported.
- Do not enable transmission of IPv6 ICMP redirect messages on RSMLT VLANs (ICMP redirect is disabled by default).

RSMLT peers not up

If, after a series of reconfigurations, RSMLT peers do not transition to the up state, use the following procedure to troubleshoot the issue. You can observe this issue on dual-stack VLANs after multiple delete and re-adds of IPv4 interfaces, or after disabling and reenabling of IPv6 forwarding or similar configurations.

Procedure

1. Display the RSMLT configuration. This command shows whether the peers are up:

```
show ip rsmlt peer
```

2. Enter VLAN Interface Configuration mode:

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

3. To recover the peers if they are down, disable and reenable RSMLT on both IST peers:

```
no ip rsmlt  
ip rsmlt
```

4. If the problem persists, boot from a saved configuration.

Example

Display the RSMLT configuration:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 1
VSP-9012:1(config-if)#show ip rsmlt peer
```

```
=====
                          Ip Rsmlt Peer Info - GlobalRouter
=====
```

VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
1	192.0.2.1	00:1f:ca:1e:d3:1e	Enable		Up	60 180
2	198.51.100.1	00:1b:ca:1d:e3:1d	Enable		Up	60 180

```
-----
```

VID	HDT REMAIN	HUT REMAIN	SMLT ID
1	60	180	10
2	60	180	10, 16

```
-----
```

VID	IPv6	MAC	ADMIN	OPER	HDTMR	HUTMR
-----	------	-----	-------	------	-------	-------

```
-----
```

VID	HDT REMAIN	HUT REMAIN	SMLT ID
-----	------------	------------	---------

```
-----VSP-9012:1(config-
if)#no ip rsmlt
VSP-9012:1(config-if)#ip rsmlt
```

Enabling trace messages for RSMLT troubleshooting

Use the following procedure to obtain additional RSMLT-related information.

Procedure

If the preceding information does not resolve the issue, you can use the following command to obtain additional RSMLT-related information:

```
trace level 15 4
```

! **Important:**

Enabling this trace on a loaded system can slow down the CPU, especially if executed through the console. Use Telnet if possible.

Troubleshooting IPv6 connectivity loss

If the switch experiences loss of IPv6 connectivity, use the following procedure to troubleshoot the issue.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Through ACLI commands, make sure the required routes are in place and the corresponding neighbor entries are resolved (that is, in REACHABLE, PROBE, DELAY or STALE state).
3. INCOMPLETE neighbor state indicates a problem if the corresponding neighbor is used by some of the IPv6 routes. This applies to neighbor entries with link-local addresses.

 **Note:**

Global addresses are not normally used as next hops. Having a global IPv6 neighbor entry as INCOMPLETE does not usually lead to a connectivity issue.

4. If the corresponding route is not in place then this is a routing issue. If the neighbor is not present or is INCOMPLETE, then further debugging is needed on the network level (that is, the state of other nodes needs to be examined).
5. Disabling and re-enabling IPv6 on the VLAN often recovers connectivity.
6. Display the RSMLT and MLT status:

```
show ip rsmlt
```

```
show mlt
```

Make sure the RSMLT peer MAC is learned and the IST state is `ist`.

Troubleshooting TACACS+

The switch supports the Terminal Access Controller Access Control System plus (TACACS+) client. TACACS+ is a remote authentication protocol that provides centralized validation of users who attempt to gain access to a router or network access server (NAS). The TACACS+ feature is disabled by default.

The current implementation of TACACS+ does not support:

- Earlier versions of TACACS
- Point-to-Point Protocol (PPP) authentication and accounting
- IPv6 addresses

TACACS+ is part of the Base Software License. For more information about licensing, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

See the following sections to troubleshoot TACACS+.

Unable to log on using Telnet or rlogin

If you cannot log on using Telnet or rlogin, perform the following steps.


Procedure

1. Check whether the TACACS+ server is available or unreachable.
2. On the TACACS+ server, check whether you configured the privilege level correctly. On successful authorization, the TACACS+ server returns an access level to the switch for the current user, which determines the user access privileges. The switch supports access levels 1 to 6 and access level 15.

The following table maps user accounts to TACACS+ privilege level.

Switch access level	TACACS+ privilege level	Description
NONE	0	If the TACACS+ server returns an access level of 0, the user is denied access. You cannot log into the device if you have an access level of 0.
READ ONLY	1	Permits you to view only configuration and status information.
LAYER 1 READ WRITE	2	Permits you to view most of the switch configuration and status information and change physical port settings.
LAYER 2 READ WRITE	3	Permits you to view and change configuration and status information for Layer 2 (bridging and switching) functions.
LAYER 3 READ WRITE	4	Permits you to view and change configuration and status information for Layer 2 and Layer 3 (routing) functions.
READ WRITE	5	Permits you to view and change configuration and status information across the VSP switch. This level does not allow you to change security and password settings.

Table continues...

Switch access level	TACACS+ privilege level	Description
READ WRITE ALL	6	Permits you to have all the rights of read-write access and the ability to change security settings, including Avaya command line interface (ACLI) and web-based management user names and passwords, and the SNMP community strings.
NONE	7 to 14	If the TACACS+ server returns an access level of 7 to 14, the user is denied access. You cannot log into the device if you have an access level of 7 to 14.
READ WRITE ALL	15	Permits you to have all the rights of read-write access and the ability to change security settings, including Avaya command line interface (ACLI) and Web-based management user names and passwords, and the SNMP community strings.  Note: Access level 15 is internally mapped to access level 6, which ensures consistency with other vendor implementations. The VSP switch does not differentiate between an access level of 6 and an access level of 15.

After you enable TACACS+ authorization, the current privilege-level to command mapping on the switch is no longer relevant because the TACACS+ server has complete responsibility for command authorization. TACACS+ authorization provides access to the system based on username, not based on privilege level.

 **Note:**

If you want to switch to a privilege level 'X' using `tacacs switch level <1-15>` command, you must create a user "\$enabX\$" on the TACACS+ server. X is the privilege level that you want to change.

3. On the TACACS+ server, check whether you configured the password and user name correctly.

4. On the TACACS+ server, check whether you configured the switch IP address in the trust list.
5. Check whether you configured the encryption key, connection mode (single connection or per-session connection), and TCP port number the same on the TACACS+ server and switch.
6. If you can log on to the switch, check whether the TACACS+ server configured on the platform has the correct IP address:

```
show tacacs
```

7. Use the output from the preceding step to verify whether the key field configured on the platform is the same as that on the TACACS+ server.
8. Also use the output from the `show tacacs` command to verify whether you configured the single connection option on the platform, and whether the TACACS+ server supports the single connection.

Example

Check whether the TACACS+ server configured on the platform has the correct IP address:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

  global enable : false

  authentication enabled for : cli

  accounting enabled for : none

  authorization : disabled

  User privilege levels set for command authorization : None

Server:

      create :

Prio  Status  Key      Port  IP address  Timeout Single Source SourceEnabled
-----
Primary NotConn ***** 3    192.0.2.254 30   true 5.5.5.5 true
Backup  NotConn ***** 47   198.51.100.1 10  false 0.0.0.0 false
```

Job aid

The following table describes the fields in the output for the `show tacacs` command.

Name	Description
Global Status	
global enable	Displays if the TACACS+ feature is enabled globally.
authentication enabled for	Displays which application is authenticated by TACACS+. The possibilities are ACLI, web, or all.

Table continues...

Name	Description
accounting enabled for	Displays if accounting is enabled. You can only enable accounting for ACLI. By default, accounting is not enabled.
authorization	Displays if authorization is enabled.
User privilege levels set for command authorization	<p>Displays the privilege levels set for command authorization. When you configure command authorization for a particular level, all commands that you execute are sent to the TACACS+ server for authorization. The device can only execute the commands the TACACS+ server authorizes.</p> <p>The user privilege levels are:</p> <ul style="list-style-type: none"> • 0: denied access • 1: read only (ro) access • 2: Layer 1 read and write (l1) access • 3: Layer 2 read and write (l2) access • 4: Layer 3 read and write (l3) access • 5: read and write (rw) access • 6: read and write all (rwa) access • 7-14: denied access • 15: read and write all (rwa) access
Server	
Prio	Displays the priority of the TACACS+ server. The switch attempts to use the primary server first, and the secondary server second.
Status	Displays the connection status between the server and the switch – connected or not connected.
Key	Displays as ***** instead of the actual key. The key is secret and is not visible.
Port	Displays the TCP port used to establish the connection to the server. The default port is 49.
IP address	Displays the IP address for the primary and secondary TACACS+ servers.
Timeout	Displays the period of time, in seconds, the switch waits for a response from the TACACS+ daemon before it times out and declares an error. The default is 10 seconds.
Single	Displays if a single open connection is maintained between the switch and TACACS+ daemon, or if the switch opens and closes the TCP connection to the

Table continues...

Name	Description
	TACACS+ daemon each time they communicate. The default is false, which means the device does not maintain the single open connection.
Source	Displays the fixed source IP address, if you configure one, for all outgoing TACACS+ packets.
SourceEnabled	Displays if the fixed source IP address is enabled for all outgoing TACACS+ packets.

Unable to log on using SSH

If you cannot log on using Secure Shell (SSH), perform the following steps.

Procedure

1. Verify that the network, the switch, and the TACACS+ server is reachable.
2. Verify whether you configured the SSH client correctly.
3. Verify whether you enabled and configured the SSH function correctly on the switch:

```
show ssh global
```

Example

Verify whether you enabled and configured SSH function correctly on the switch:

```
Switch:1>enable
Switch:1#show ssh global

Total Active Sessions : 0
  version              : v2only
  port                 : 22
  max-sessions         : 4
  timeout              : 60
  action rsa-keygen    : rsa-keysize 1024
  action dsa-keygen    : dsa-keysize 1024
  rsa-auth             : true
  dsa-auth             : true
  pass-auth            : true
  enable               : false
```

Job Aid

The following table describes the fields in the output for the `show ssh global` command.

Parameter	Description
Total active sessions	Specifies the number of active SSH sessions underway.
version	Specifies if SSH is version 1 or version 2. The default is v2. Avaya recommends you configure the version to v2 only.

Table continues...

Parameter	Description
port	Specifies the SSH connection port. The default is 22. You cannot configure the following TCP ports as SSH connection ports: 0 to 1024 (except port 22), 1100, 4095, 5000, 5111, 6000, or 999.
max-sessions	Specifies the maximum number of SSH sessions allowed. The default is 4.
timeout	Specifies the SSH connection authentication timeout in seconds. The default is 60 seconds.
action rsa-keygen	Specifies the SSH RSA key size.
action dsa-keygen	Specifies the SSH DSA key size.
rsa-auth	Specifies if RSA authentication is enabled or disabled. The default is enabled.
dsa-auth	Specifies if DSA authentication is enabled or disabled. The default is enabled.
pass-auth	Specifies if password authentication is enabled or disabled. The default is enabled.
enable	Specifies if SSH secure mode is enabled. False is disabled. Secure is enabled.

Unable to log on by any means (Telnet, rlogin, or SSH)

If you cannot log on by any means, perform the following steps.

Procedure

1. Check whether the TACACS+ server runs properly and try to restart the TACACS+ server.
2. Check whether you enabled both TACACS+ and RADIUS on the switch.

```
show radius
```

```
show tacacs
```

If TACACS+ fails, RADIUS can take over the authentication, authorization, and accounting (AAA) process.

3. Check whether you configured the TACACS+ server to unencrypted mode, as the switch always sends encrypted TACACS+ messages.
4. Check whether you configured the switch properly. In particular, check the IP address and key.

```
show tacacs
```

5. Check whether you configured the encryption key, connection mode (single connection or per-session connection), and TCP port number the same on the TACACS+ server and switch.

6. If the server connects directly, check whether the administrative and operation status of the port is up:

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

7. If the server is connected in a network, check whether the switch has a route configured to the server network:

```
show ip route
```

Example

Check whether you enabled both TACACS+ and RADIUS on the switch:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

  global enable : false

  authentication enabled for : cli

  accounting enabled for : none

  authorization : disabled

  User privilege levels set for command authorization : None

Server:

          create :

Prio   Status  Key      Port  IP address  Timeout  Single Source  SourceEnabled
Primary NotConn *****  3     192.0.2.254   30    true 5.5.5.5  true
Backup  NotConn *****  47    198.51.100.1  10    false 0.0.0.0 false

Switch:1(config)#show radius
  acct-attribute-value : 193
  acct-enable : false
  acct-include-cli-commands : false
  access-priority-attribute : 192
  auth-info-attr-value : 91
  command-access-attribute : 194
  cli-commands-attribute : 195
  cli-cmd-count : 40
  cli-profile-enable : false
  enable : false
  igap-passwd-attr : standard
  igap-timeout-log-fsize : 512
  maxserver : 10
  mcast-addr-attr-value : 90
  sourceip-flag : false
```

Check whether the administrative and operation status of the port is up:

```
Switch:1#show interface gigabitethernet 4/2

=====
                          Port Interface
=====
PORT      LINK  PORT  PHYSICAL  STATUS
NUM  INDEX DESCRIPTION  TRAP  LOCK  MTU  ADDRESS  ADMIN  OPERATE
=====
```

```

4/2 257 1000BaseTX true false 1950 00:24:7f:a1:70:61 up up
=====
Port Name
=====
PORT OPERATE OPERATE OPERATE
NUM NAME DESCRIPTION STATUS DUPLX SPEED VL
AN
-----
4/2 1000BaseTX up full 1000 Ta
gged
=====
Port Config
=====
PORT DIFF-SERV QOS MLT VENDOR
--More-- (q = quit)

```

Check whether the switch has a route configured to the server network:

```

Switch:1(config)#show ip route
=====
IP Route - GlobalRouter
=====
NH
INTER DST MASK NEXT VRF/ISID COST FACE PROT AGE
TYPE PRF
-----
198.51.100.1 255.255.255.255 192.0.2.65 GlobalRouter 1 100 OSPF 0
IB 125
198.51.100.5 255.255.255.255 192.0.2.5 - 1 0 LOC 0
DB 0
198.51.100.13 255.255.255.255 VSP13 GlobalRouter 10 1000 ISIS 0
IBS 7
198.51.100.200 255.255.255.255 VSP200 GlobalRouter 10 1000 ISIS 0
IBS 7
4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
-----
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

Job aid

The following table describes the fields in the output for the `show radius` command.

Parameter	Description
acct-attribute-value	Specifies the accounting attribute value.
acct-enable	Specifies if the accounting attribute is enabled.

Table continues...

Parameter	Description
acct-include-cli-commands	Specifies if the accounting attribute includes ACLI commands. The default is false.
access-priority-attribute	Specifies the value of the access priority attribute. The default is 192.
auth-info-attr-value	Specifies the value of the authentication information attribute. The default is 91.
command-access-attribute	Specifies the value of the command access attribute. The default is 194.
cli-commands-attribute	Specifies the value of the ACLI commands attribute. The default is 195.
cli-cmd-count	Specifies how many ACLI commands before the system sends a RADIUS accounting interim request. The default is 40.
cli-profile-enable	Specifies if RADIUS ACLI profiling is enabled. ACLI profiling grants or denies access to users being authenticated by way of the RADIUS server. You can add a set of ACLI commands to the configuration on the RADIUS server, and you can specify the command-access mode for these commands. The default is false.
enable	Specifies if RADIUS authentication is globally enabled on the switch.
igap-passwd-attr	Specifies the IGMP for user Authentication Protocol (IGAP) password attribute.
igap-timeout-log-fsize	Specifies the IGMP for user Authentication Protocol (IGAP) timeout log file size.
maxserver	Specifies the maximum number of servers allowed for the device. The default is 10.
mcast-addr-attr-value	Specifies the value of the multicast address attribute. The default is 90.
sourceip-flag	<p>Specifies if the switch can use a configured source IP address. If the outgoing interface on the switch fails, a different source IP address is used, which requires that you make configuration changes to define the new RADIUS client on the RADIUS server. To simplify RADIUS server configuration, you can configure the switch to use a circuitless IP (CLIP) address as the source IP and NAS IP address when transmitting RADIUS packets.</p> <p>By default, the switch uses the IP address of the outgoing interface as the source IP, and the NAS IP address for RADIUS packets that it transmits.</p>

Administrator unable to obtain accounting information from the TACACS+ server

If the administrator is unable to obtain accounting information from the TACACS+ server, perform the following steps.

Procedure

1. Check whether you enabled accounting on the switch:


```
show tacacs
```
2. Check whether you enabled accounting on the TACACS+ server.

Example

Check whether accounting is enabled on the switch:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

  global enable : false

  authentication enabled for : cli

  accounting enabled for : none

  authorization : disabled

  User privilege levels set for command authorization : None

Server:

          create :

Prio   Status  Key      Port  IP address  Timeout Single Source SourceEnabled
Primary NotConn *****  3     192.0.2.254   30    true 5.5.5.5  true
Backup  NotConn *****  47    198.51.100.1  10    false 0.0.0.0 false
```

Trap server cannot receive trap packets from the VSP device

If the trap server cannot receive trap packets from the switch, perform the following steps.

Procedure

1. Check whether you configured the trap server correctly on the switch:


```
show snmp-server host
```
2. Check whether a firewall exists between the switch and the trap server.

Example

Check whether you configured the trap server correctly on the switch:

```
Switch:1>enable
Switch:1#show snmp-server host
```

```
=====
Notify Configuration
=====
```

Notify Name	Tag	Type
Inform	informTag	inform
Trap	trapTag	trap

```
=====
Notify Profile Configuration
=====
```

Params Name	Profile Name
AuthNoPriv-md5	profile2
AuthPriv-md5	profile3
NoAuthNoPriv-md5	profile1

```
=====
Target Address Configuration
=====
```

Target Name	TDomain	TAddress	TMask
4c20cc369925edbd1fe3cf8e2584c498	ipv4	47.17.142.155:162	
55fca382ffba169e986783bbbdedc334	ipv4	47.17.143.57:162	

```
=====
Target Address Configuration
=====
```

Target Name Params	Timeout MMS	Retry	TagList
4c20cc369925edbd1fe3cf8e2584c498	1500	3	trapTag
4c20cc369925edbd1fe3cf8e2584c498	484		
55fca382ffba169e986783bbbdedc334	1500	3	trapTag
55fca382ffba169e986783bbbdedc334	484		

```
=====
Target Params Configuration
=====
```

Target Name Level	MP Model	Security Name	Sec
4c20cc369925edbd1fe3cf8e2584c498	snmpv1	readview	noAu
thNoPriv			
55fca382ffba169e986783bbbdedc334	snmpv2c	secret	noAu
thNoPriv			
TparamV1	snmpv1	readview	noAu
thNoPriv			
TparamV2	snmpv2c	readview	noAu
thNoPriv			

Troubleshooting TACACS+ problems

Use the `trace level` command to check traps and log files to see any TACACS+ failure. If TACACS+ experiences failure conditions, the TACACS+ module sends SNMP traps to notify the user. The TACACS+ module also logs the failure information into the system log file.

About this task

Caution:

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation. If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends that you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Configure the trace level for the TACACS+ module:

```
trace level 109 <1-4>
```

The TACACS+ module ID is 109.

3. Stop trace:

```
trace shutdown
```

4. View the trace results on screen:

```
trace screen enable
```

5. View trace saved to a file:

```
show trace file [tail]
```


6. Save the trace file to the Compact Flash card for retrieval:

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is `systrace.txt`. By default, the system saves the file to the external flash.

Variable definitions

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

Variable	Value
level [<code><0-192></code>][<code><1-4></code>]	Starts the trace by specifying the module ID and level. <code><0-192></code> specifies the module ID. Module ID 23 represents the IGMP module <code><1-4></code> specifies the trace level: <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable enable}	Enables or disables the display of trace output to the screen.  Important: Avaya recommends you avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

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

Variable	Value
file [tail]	Displays the trace results saved to a file.
level	Displays the current trace level for all modules.
modid-list	Specifies the module ID list.

Troubleshooting client registration

About this task

Perform this procedure if a client is not registered by the switch.

Procedure

1. Enable auto-negotiation on the client port.
2. Disable and enable the port.

Chapter 12: Unicast routing troubleshooting

Use this section to troubleshoot Layer 3 unicast routing problems.

Using BGP debugging commands

Use global and peer debug commands to display specific debug messages for the global and peer Border Gateway Protocol (BGP) configuration, including the BGP neighbors.

You can use these commands to troubleshoot the BGP configuration.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
configure terminal
router bgp
```

2. Show specific debug messages for the global BGP configuration:

```
global-debug mask WORD<1-100>
```

3. Display specific debug messages for the global BGP neighbors:

```
neighbor-debug-all mask WORD<1-100>
```

4. Display specific debug messages for BGP peers or peer groups:

```
neighbor <nbr_ipaddr|peer-group-name> neighbor-debug-mask
WORD<1-100>
```

5. Display debug messages on the console:

```
debug-screen <on|off>
```

Example

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

```
Switch:1(config)# router bgp
```

Display the global debug messages for error and packet:

```
Switch:1(router-bgp)#global-debug mask error,packet
```

End (disable) the display of global debug messages:

```
Switch:1(router-bgp)#global-debug mask none
```

Display specific debug messages for the global BGP neighbors:

```
Switch:1(router-bgp)#neighbor-debug-all mask packet,event
```

Display specific debug messages for BGP peers or peer groups:

```
Switch:1(router-bgp)#neighbor 45.17.10.23 neighbor-debug-mask event,trace
```

Display debug messages on the console:

```
Switch:1(router-bgp)#debug-screen on
```

Variable definitions

Use the data in the following table to use the **global-debug mask** and **neighbor-debug-all mask** commands.

Variable	Value
<i>WORD</i> <1-100>	Specifies one or more mask choices that you enter, separated by commas with no space between choices. For example: [<i><mask></i> , <i><mask></i> , <i><mask></i> ...]. Options include: none, all, error, packet, event, trace, warning, state, init, filter, update.

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

Variable	Value
<nbr_ipaddr peer-group-name>	Specifies the IP address or the group name of the peer.
<i>WORD</i> <1-100>	Specifies one or more mask choices that you enter, separated by commas with no space between choices. For example: [<i><mask></i> , <i><mask></i> , <i><mask></i> ...]. Options include: none, all, error, packet, event, trace, warning, state, init, filter, update.

Job aid

Use debug command values to control debug messages for global BGP message types, and for message types associated with a specified BGP peer or peer group. The following table identifies mask categories and messages.

Table 42: Mask categories and messages

Mask category	Message
none	None disables the display of all debug messages.
all	All configures the device to show all categories of debug messages.
error	Error configures the device to show error debug messages.
packet	Packet configures the device to show packet debug messages.
event	Event configures the device to show event debug messages.
warning	Warning configures the device to show warning debug messages.
init	Init configures the device to show initialization debug messages.
filter	Filter configures the device to show filter-related debug messages.
update	Update configures the device to show update-related debug messages.

Troubleshooting licensed routing protocols

About this task

Many routing protocols require a license for operation. Perform this procedure if a licensed protocol does not operate.

For more information about how to install or transfer licenses, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

Procedure

1. Verify that the license is the correct type.
2. Verify that you installed the license properly.

Example

Verify that the license is the correct type:

```
VSP-9012:1#show license
```

```

License file name      : /intflash/premiersitelicense.dat
License Type          : PREMIER
MD5 of Key             : 67d9b8d4 e58172cf 91a3a4c2 5f03c00a
MD5 of File           : 185f5e5c fea563d0 dd1a777c 8d54208c
Generation Time       : 2010/04/12 11:18:08
Expiration Time       :
Base Mac Addr         : 00:24:7f:9f:60:00

```

```
flags          : 0x00000001 SINGLE
memo          :
```

Job aid

Base License

The Base License activates the features not included in either the Advanced or Premier Licenses.

- Terminal Access Controller Access Control System Plus (TACACS+)

Advanced License

The Advanced License activates the following features in addition to the Base License:

- Border Gateway Protocol version 4 (BGP4) for 16 peers or 64,000 routes
- Packet Capture function (PCAP)
- Layer 3 mirroring
- IPv6 routing

Premier License

The Premier License activates the following features in addition to the Advanced License:

- Virtual Routing and Forwarding (VRF)
- IP routes:
 - If you configure VSP 9000 for first generation I/O modules, the device scales to 400,000 IPv4 FIB routes.
 - If you configure VSP 9000 for second generation I/O modules, the device scales to 1,000,000 IPv4 FIB routes.
- 256 BGP peers
- Lossless Ethernet
- Shortest Path Bridging MAC (SPBM) – L3 VSNs, Inter-VSN routing, IP Shortcuts
- IP multicast over Shortest Path Bridging MAC (SPBM)

The Premier License activates all licensed features on Virtual Services Platform 9000.

Viewing OSPF errors

Check Open Shortest Path First (OSPF) errors for administrative and troubleshooting purposes.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display information about OSPF errors:

```
show ip ospf port-error [port {slot/port[-slot/port][,...]}] [vrf
WORD<1-16>] [vrfs WORD<0-512>]
```

Example

Display information about OSPF errors:

```
VSP-9012:1>enable
VSP-9012:1(config)#show ip ospf port-error
```

Variable definitions

Use the data in the following table to use the `show ip ospf port-error` command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (3/1), a range of slots and ports (3/2-3/4), or a series of slots and ports (3/2,5/3,6/2).
vrf WORD<1-16>	Specifies a VRF by name.
vrfs WORD<0-512>	Specifies a range of VRF IDs.

Job aid

The following table explains the fields in the `show ip ospf port-error` command output.

Table 43: OSPF port error field descriptions

Field	Description
PORT NUM	Indicates the port number.
VERSION MISMATCH	Indicates the number of version mismatches this interface receives.
AREA MISMATCH	Indicates the number of area mismatches this interface receives.
AUTHYPEMISMATCH	Indicates the number of authentication type mismatches this interface receives.
AUTH FAILURES	Indicates the number of authentication failures.
NET_MASK MISMATCH	Indicates the number of network mask mismatches this interface receives.
HELLOINT MISMATCH	Indicates the number of hello interval mismatches this interface receives.

Table continues...

Field	Description
DEADINT MISMATCH	Indicates the number of dead interval mismatches this interface receives.
OPTION MISMATCH	Indicates the number of options mismatches this interface receives.

Viewing OSPF neighbor state problems

About this task

View the status of all the OSPF neighbors and their current adjacency state to determine if problems occurred during the device initial startup sequence.

Problems with OSPF occur most often during the initial startup, when the device cannot form adjacencies with other devices, and the state is stuck in the Init or ExStart/Exchange state.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the current state of all OSPF neighbors and their current state of adjacency:

```
show ip ospf neighbor
```

Example

View the current state of all OSPF neighbors and their current state of adjacency:

```
VSP-9012:1>enable
VSP-9012:1#show ip ospf neighbor
```

```
=====
                        OSPF Neighbors - GlobalRouter
=====
INTERFACE      NBRROUTERID    NBRIPADDR      PRIO_STATE    RTXQLEN  PERM  TTL
-----
42.1.1.33      198.95.65.0   42.1.1.34      1    Full    0     Dyn    40
=====
```

Job aid

At initial startup, devices transmit hello packets in an attempt to find other OSPF devices with which to form adjacencies. After the device receives the hello packets, it performs an initialization process, which causes the device to transition through various states before it establishes the adjacency.

The following table describes the various device states during adjacency formation.

Table 44: Device states during OSPF adjacency formation

Step	State	Description
1	Down	Indicates that a neighbor was configured manually, but the device did not receive information from the other device. This state can occur only on nonbroadcast multiaccess interfaces.
2	Attempt	Indicates, on a nonbroadcast multiaccess interface, that the device attempts to send unicast hellos to configured interfaces.
3	Init	Indicates that the device received a general hello packet (without the router ID) from another device.
4	2-Way	Indicates that the device received a hello packet directed to it from another device. (The hello contains the router ID.)
5	ExStart	Indicates the start of the master and backup election process.
6	Exchange	Indicates the link state database is exchanged.
7	Loading	Indicates the processing state of the link state database for input into the routing table. The device can request link state advertisements for missing or corrupt routes.
8	Full	Indicates the normal full adjacency state.

Troubleshooting OSPF Init state problems

About this task

A device can become stuck in the Init state and not form an adjacency. Several possible causes for this type of problem exist:

- Access lists implemented on routers.
- Authentication mismatch or configuration problem.

Problems arise if a mismatch exists in authentication keys, or if both sides are not configured for authentication.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Configure the trace level for the OSPF module to terse:


```
trace level 6 2
```
3. View the trace information on screen:


```
trace screen enable
```
4. Verify if the path is not reachable due to access lists implemented on the device.

5. Ensure the multicast address of 224.0.0.5 can traverse the link. If multicast traffic is blocked, you must configure the Avaya Virtual Services Platform 9000 for OSPF nonbroadcast multiaccess (NBMA) instead of broadcast.

Example

Configure the trace level for the OSPF module to terse and view the trace information on screen:

```
VSP-9012:1>enable
VSP-9012:1#trace level 6 2
VSP-9012:1#trace screen enable
```

Troubleshooting OSPF ExStart/Exchange problems

About this task

Although both devices can recognize each other and move beyond 2-way state, the devices can become stuck in the ExStart/Exchange state. A mismatch in maximum transmission unit (MTU) sizes between the devices usually causes this type of problem. For example, one device can use a high MTU size and the default value on the other device is a smaller value. Depending on the size of the link state database (LSDB), the device with the smaller value cannot process the larger packets and remains in ExStart/Exchange state. This problem is usually encountered during interoperations in networks with other vendor devices.

In Virtual Services Platform 9000, the supported MTU size for OSPF is 1500 bytes by default. Incoming OSPF database description (DD) packets are dropped if their MTU size is greater than 1500 bytes.

If you configure the device to ignore the MTU size, the device does not perform the MTU check on the incoming OSPF DD packet. Virtual Services Platform 9000 automatically checks for OSPF MTU mismatches.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

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

2. View the OSPF packets:

```
trace level 6 2
```

3. Ensure that the MTU size value for both devices match.
4. Configure the interface to accept OSPF DD packets with a different MTU size:

```
ip ospf mtu-ignore enable
```


Example

View the OSPF packets. Configure the interface to accept OSPF DD packets with a different MTU size.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 100
VSP-9012:1(config-if)#trace level 6 2
VSP-9012:1(config-if)#ip ospf mtu-ignore enable
```

Chapter 13: Multicast troubleshooting

Use the following information to troubleshoot multicast features and multicast routing.

Multicast feature troubleshooting

Use the information in this section to troubleshoot multicast feature problems.

Troubleshooting IGMP Layer 2 Querier

The following sections provide troubleshooting information for the IGMP Layer 2 Querier feature.

Querier not elected

If a Querier is not elected, use the following procedure to troubleshoot the issue.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. As the IGMP Layer 2 Querier is based on IGMP snoop, check whether IGMP snoop is enabled on the VLAN:

```
show ip igmp interface vlan
```

If IGMP snoop is disabled, the Layer 2 Querier cannot work until IGMP snoop and IGMP Layer 2 Querier are reenabled.

Example

Check whether IGMP snoop is enabled on the VLAN:

```
Switch:1>enable  
Switch:1#show ip igmp interface vlan
```

```
=====
                        Vlan Ip Igmp
=====
VLAN QUERY QUERY ROBUST VERSION LAST  PROXY  SNOOP  SSM    FAST  FAST
ID   INTVL MAX    RESP          MEMB  SNOOP  ENABLE SNOOP  LEAVE LEAVE
                        QUERY ENABLE      ENABLE ENABLE PORTS
=====
```

```

1      125    100    2      2      10      false  false  false  false
2      125    100    2      2      10      false  false  false  false
3      125    100    2      2      10      false  false  false  false
4      125    100    2      2      10      false  false  false  false
5      125    100    2      2      10      false  false  false  false
10     125    100    2      2      10      false  false  false  false
100    125    100    2      2      10      false  false  false  false
200    125    100    2      2      10      false  false  false  false
300    125    100    2      2      10      false  false  false  false
444    125    100    2      2      10      false  false  false  false

```

All 10 out of 10 Total Num of Icmp entries displayed

VLAN ID	SNOOP QUERIER	SNOOP QUERIER ADDRESS	DYNAMIC DOWNGRADE VERSION	COMPATIBILITY MODE	EXPLICIT HOST TRACKING
1	false	0.0.0.0	enable	disable	disable
2	false	0.0.0.0	enable	disable	disable
3	false	0.0.0.0	enable	disable	disable
4	false	0.0.0.0	enable	disable	disable
5	false	0.0.0.0	enable	disable	disable
10	false	0.0.0.0	enable	disable	disable
100	false	0.0.0.0	enable	disable	disable
200	false	0.0.0.0	enable	disable	disable
300	false	0.0.0.0	enable	disable	disable
444	false	0.0.0.0	enable	disable	disable

All 10 out of 10 Total Num of Icmp entries displayed

Job aid

The following table describes the fields in the output for the `show ip igmp interface vlan` command.

* Note:

The following table shows the field descriptions for this command if you use the optional parameter `vlan`. If you do not the output is different.

Field	Description
VLAN ID	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that runs on this interface. This object configures a router capable of

Table continues...

Field	Description
	running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.
PROXY SNOOP ENABLE	Indicates if proxy snoop is enabled on the interface.
SNOOP ENABLE	Indicates if snoop is enabled on the interface.
SSM SNOOP ENABLE	Indicates if SSM snoop is enabled on the interface.
FAST LEAVE ENABLE	Indicates if fast leave mode is enabled on the interface.
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave.
VLAN ID	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
SNOOP QUERIER ENABLE	Specifies whether the snoop querier is enabled.
SNOOP QUERIER ADDRESS	Specifies the pseudo address of the IGMP snoop querier.
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates whether compatibility mode is enabled.
EXPLICIT HOST TRACKING	Specifies whether the IGMP protocol version 3 is enabled to track hosts for each channel or groups.

Enabling trace messages for IGMP Layer 2 querier troubleshooting

If the preceding information does not address your issue, you can also use the following trace command to view additional information related to Layer 2 querier.

Caution:

Using the trace tool inappropriately can cause primary CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation. If you use trace level 3 (verbose) or trace level 4 (very verbose), Avaya recommends that you do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Use the following trace command to begin the trace operation for additional information related to Layer 2 querier:

```
trace level 23 <1-4>
```

3. Stop tracing:

```
trace shutdown
```

4. View the trace results:

```
trace screen enable
```

5. View trace saved to a file:

```
show trace file [tail]
```

6. Save the trace file to the Compact Flash card for retrieval:

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

Variable definitions

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


Variable	Value
level [<0-207>][<1-4>]	<p>Starts the trace by specifying the module ID and level. <0-207> specifies the module ID. Module ID 23 represents the IGMP module.</p> <p><1-4> specifies the trace level:</p> <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
shutdown	Stops the trace operation.
screen {disable enable}	<p>Enables or disables the display of trace output to the screen.</p> <p> Important:</p> <p>Avaya recommends you avoid using the screen to view commands if you use trace level 3</p>


Table continues...

Variable	Value
	(verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

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

Variable	Value
file [tail]	Displays the trace results saved to a file.
level	Displays the current trace level for all modules.
modid-list	Specifies the module ID list.

Use the data in the following table to use the **save trace** command.

Variable	Value
file <i>WORD</i> <1–99>	<p>Specifies the file name in one of the following formats:</p> <ul style="list-style-type: none"> • a.b.c.d: <file> • x:x:x:x:x:x:x: <file> • /intflash/ <file> • /extflash/ <file> • /mnt/intflash/ <file> • /mnt/extflash/ <file> • /usb/<file> <p>/mnt/intflash/ is the internal flash of the second CP module (the one to which you are not connected.)</p> <p>/mnt/extflash/ is the external flash of the second CP module (the one to which you are not connected.)</p> <p><i>WORD</i><1–99> is a string of 1–99 characters.</p> <p> Note:</p> <p>If you do not specify a file name, the file name is <code>systrace.txt</code>. By default, the system saves the file to the external flash.</p>

Troubleshooting IGMPv3 backwards compatibility

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv2 and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message. If it is a v2 message, IGMP snoop processes handle the message.

To troubleshoot issues with the IGMPv3 backwards compatibility feature, perform the following procedure.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Verify that the SSM static channel is configured for the v1/v2 joins received. Display the configured SSM static channels:

```
show ip igmp ssm-map
```

3. Verify that the SSM group range is configured for the v1/v2 joins received. Display the configured SSM group range:

```
show ip igmp ssm
```

Example

Display the configured SSM static channels and display the configured SSM group range:

```
Switch:>enable
Switch:1#show ip igmp ssm-map
```

Igmp Ssm Channel				
GROUP	SOURCE	MODE	ACTIVE	STATUS
233.252.0.1	192.0.2.200	dynamic	false	enabled
233.252.0.2	192.0.2.200	dynamic	false	enabled
233.252.0.3	192.0.2.200	dynamic	false	enabled
233.252.0.4	192.0.2.200	dynamic	false	enabled
233.252.0.5	192.0.2.200	dynamic	false	enabled
233.252.0.6	192.0.2.200	dynamic	false	enabled
233.252.0.7	192.0.2.200	dynamic	false	enabled
233.252.0.8	192.0.2.200	dynamic	false	enabled
233.252.0.9	192.0.2.200	dynamic	false	enabled
233.252.0.10	192.0.2.200	dynamic	false	enabled

```
10 out of 10 entries displayed
Switch:1(config)#show ip igmp ssm
```

Igmp Ssm Global - GlobalRouter	
DYNAMIC LEARNING	SSM GROUP RANGE
enable	233.252.0.0/255.0.0.0

Job aid

The following table shows the field descriptions for the `show ip igmp ssm-map` command.

Table 45: show ip igmp ssm-map command

Field	Description
GROUP	Indicates the IP multicast group address that uses the default range of 232/8.
SOURCE	Indicates the IP address of the source that sends traffic to the group source.
MODE	Indicates that the entry is a statically configured entry (static) or a dynamically learned entry from IGMPv3 (dynamic).
ACTIVE	Indicates the activity on the corresponding source and group. If the source is active and traffic is flowing to the switch, this status is active; otherwise, it is nonactive.
STATUS	Indicates the administrative state and whether to use the entry. If the status is enabled (default), the entry is used. If the status is disabled, the entry is not used but is saved for future use.

The following table shows the field descriptions for the `show ip igmp ssm` command.

Table 46: show ip igmp ssm command

Field	Description
DYNAMIC LEARNING	Indicates whether dynamic learning is enabled at a global level.
SSM GROUP RANGE	Indicates the IP address range for the SSM group.

Multicast routing troubleshooting using ACLI

Use the information in this section to help you troubleshoot multicast routing problems.

Viewing IGMP interface information

Perform this procedure to view the IGMP interface table.

About this task

If an interface does not use an IP address, it does not appear in the IGMP table. One exception is an IGMP snooping interface, which does not require an interface IP address.

If an interface uses an IP address, but neither IGMP snoop or PIM is enabled, the interface appears as inactive in the Status field.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. View IGMP interfaces:


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

Example

View IGMP interfaces:

```
Switch:1#show ip igmp interface
```

```
=====
                        IgmP Interface - GlobalRouter
=====
IF          QUERY   OPER   QUERY   WRONG   LASTMEM
  INTVL  STATUS  VERS.  VERS  QUERIER  MAXRSPT  QUERY  JOINS  ROBUST  QUERY  MODE
-----
P4/11    125    inact   3      3    0.0.0.0   100    0     0     2     10  routed-spb
P4/12    125    inact   2      2    0.0.0.0   100    0     0     2     10
P4/23    125    inact   2      2    0.0.0.0   100    0     0     2     10
V2       125    inact   2      2    0.0.0.0   100    0     0     2     10
V100     125    activ   2      2    0.0.0.0   100    0     0     2     10  routed-spb
```

```
5 out of 5 entries displayed
```

Variable definitions

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

Variable	Value
gigabitethernet {slot/port[-slot/port] [,...]}	<p>Optionally, identifies the slot and port in one of the following formats: a single slot and port, a range of slots and ports, or a series of slots and ports.</p> <p>If you do not specify a slot and port, the command output includes all IGMP interfaces.</p>
vlan <1-4084>	<p>Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>If you do not specify a VLAN ID, the command output includes all IGMP interfaces.</p>
vrf WORD <1-16>	<p>Optionally, identifies the VRF name. If you do not specify a VRF name, the results display information for the Global Router. If you specify a VRF name, the results display information only for the VRF you specify.</p>
vrfids WORD <0-512>	<p>Optionally, identifies the VRF ID. If you do not specify a range of VRF IDs, the results display information for the Global Router. If you specify a VRF ID or range of VRF IDs, the results display information only for the VRF you specify.</p>

Job aid

The following table shows the field descriptions for the command output if you do not use the optional parameters.

Table 47: show ip igmp interface command output without parameters

Field	Description
IF	Indicates the interface where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
STATUS	Indicates the activation of a row, which activates IGMP on the interface. The destruction of a row disables IGMP on the interface.
VERS.	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface attaches.
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received whose IGMP version does not match the interface version. You must configure all routers on a LAN to run the same version of IGMP. If queries are received with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times this interface added a group membership.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
LASTMEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.

The following table shows the field descriptions for the command output if you use the interface parameters.

Table 48: show ip igmp interface command output with interface parameters

Field	Description
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.
PROXY SNOOP ENABLE	Indicates if proxy snoop is enabled on the interface.
SNOOP ENABLE	Indicates if snoop is enabled on the interface.
SSM SNOOP ENABLE	Indicates if SSM snoop is enabled on the interface.
FAST LEAVE ENABLE	Indicates if fast leave mode is enabled on the interface.
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave.
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates if compatibility mode is enabled.
EXPLICIT HOST TRACKING	Indicates if explicit host tracking is enabled for IGMPv3. Explicit host tracking enables the IGMP to track all source and group members.

Viewing multicast group trace information for IGMP snoop

About this task

Multicast group trace tracks the data flow path of the multicast streams.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the multicast group trace for an IGMP snoop-enabled interface:

```
show ip igmp snoop-trace [source {A.B.C.D}] [group {A.B.C.D}]
```

Example

Display the multicast group trace for an IGMP snoop-enabled interface:

```
Switch:1>enable
Switch:1#show ip igmp snoop-trace
```

```
=====
IGMP Snoop Trace - GlobalRouter
=====
```

GROUP ADDRESS	SOURCE ADDRESS	IN VLAN	IN PORT	OUT VLAN	OUT PORT	TYPE
233.252.0.1	192.0.2.6	500	spb	500	9/5	NETWORK
233.252.0.100	192.0.2.7	500	spb	500	10/10	NETWORK

Variable definitions

Use the data in the following table to use the `show ip igmp snoop-trace` command.

Table 49: Variable definitions

Variable	Value
group {A.B.C.D}	Specifies the group IP address in the format a.b.c.d.
source {A.B.C.D}	Specifies the source IP address in the format a.b.c.d.

Job aid

The following table shows the field descriptions for the `show ip igmp snoop-trace` command.

Table 50: show ip igmp snoop-trace field descriptions

Field	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.

Table continues...

Field	Description
OUT PORT	Indicates the outgoing port number.
TYPE	Indicates where the stream is learned. If ACCESS displays, then the stream is learned on UNI ports. If NETWORK displays, then the stream is learned on the SPBM cloud.

Viewing IGMP group information

View information about IGMP groups to see the current group operation on the switch.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View IGMP group information:

```
show ip igmp group group <A.B.C.D> detail [port {slot/port[-slot/port]}[,...]] [vlan <1-4084>] [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

```
show ip igmp group group <A.B.C.D> tracked-members [member-subnet <A.B.C.D./X>] [port {slot/port[-slot/port]}[,...]] [source-subnet <A.B.C.D./X>] [vlan <1-4084>] [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

Example

View IGMP group information:

```
Switch:1>enable
Switch:1#show ip igmp group group 233.252.0.100
```

```
=====
                          Igmp Group
=====
GRPADDR      INPORT      MEMBER      EXPIRATION  TYPE
-----
233.252.0.100  V20-5/19    192.0.2.3    138         Dynamic
233.252.0.100  V20-5/19    192.0.2.4    176         Dynamic

2 out of 2 group Receivers displayed

Total number of unique groups 1
```

Variable definitions

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

Variable	Value
count	Displays the number of entries in the IGMP group.
group <A.B.C.D>	Specifies the address of the IGMP group.

Table continues...

Variable	Value
member-subnet {default <A.B.C.D>}}	Specifies the IP address and mask of the IGMP member.
vrf WORD<1–16>	Specifies the VRF name.
vrfids WORD<0–512>	Specifies the VRF ID.

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

Variable	Value
detail [port {slot/port[-slot/port][,...]} vlan <1-4084> vrfWORD <1–16> vrfidsWORD <0–255>]	<p>Use the detail parameter to show IGMPv3–specific data.</p> <p>For data related to a specific interface use the following:</p> <ul style="list-style-type: none"> port {slot/port[-slot/port][,...]} — Specifies the port list. vlan <1-4084>— Specifies the VLAN. <p>Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <ul style="list-style-type: none"> vrf WORD<1–16> — Specifies the VRF name. vrfids WORD<0–255> — Specifies the VRF ID.
tracked-members	Use the tracked-members parameter to view all the tracked members for a specific group.
vrf WORD<1–16>	Specifies the VRF name.
vrfids WORD<0–512>	Specifies the VRF ID.

Use the data in the following table to use the **show ip igmp group group <A.B.C.D> tracked-members** command.

Variable	Value
member-subnet {default <A.B.C.D>}}	Specifies the IP address and mask of the IGMP member.
port {slot/port[-slot/port][,...]}	Specifies the port list.
source-subnet <A.B.C.D/X>	Specifies the source IP address and the subnet mask.
vlan <1-4084>	Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf WORD<1–16>	Specifies the VRF name.
vrfids WORD<0–512>	Specifies the VRF ID.

Job aid

The following table shows the field descriptions for the `show ip igmp group group` command output.

Table 51: show ip igmp group group command output

Field	Description
GRPADDR	Shows the multicast group address (Class D). A group address can be the same for many incoming ports.
INPORT	Shows the port that receives the group membership report.
MEMBER	Shows the IP address of the host that issues the membership report to this group.
EXPIRATION	Shows the time left before the group report expires on this port. This variable is updated after the port receives a group report.
TYPE	Indicates the group type.

Showing the hardware resource usage

About this task

The switch can query the number of ingress and egress IP multicast streams traversing the switch. After you configure the thresholds for ingress and egress records, if the record-usage goes beyond the threshold, the device notifies you by way of a trap on the console, logged message, or both.

If you do not configure the thresholds, ACLI displays only the ingress and egress records currently in use.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Show the hardware resource usage:

```
show ip mroute hw-resource-usage
```

Example

Show the hardware resource usage:

```
Switch:1>show ip mroute hw-resource-usage
=====
                          Multicast Hardware Resource Usage
=====
EGRESS      INGRESS      EGRESS      INGRESS      LOG MSG      SEND TRAP      SEND TRAP
REC IN-USE  REC IN-USE  THRESHOLD   THRESHOLD   ONLY         ONLY          AND LOG
```

```
-----
0          0          0          0          false    false    false
```

Job aid

The following table shows the field descriptions for the `show ip mroute hw-resource-usage` command.

Table 52: show ip mroute-hw resource usage field descriptions

Field	Description
EGRESS REC IN-USE	Indicates the number of egress records (peps) traversing the switch that are in use.
INGRESS REC IN-USE	Indicates the number of source and group records traversing the switch that are in use.
EGRESS THRESHOLD	Indicates the egress records threshold.
INGRESS THRESHOLD	Indicates the source and group records threshold.
LOG MSG ONLY	Indicates the status of logging messages only.
SEND TRAP ONLY	Indicates the status of sending traps only.
SEND TRAP AND LOG	Indicates the status of both sending traps and logging messages.

Using PIM debugging commands

About this task

Use Protocol Independent Multicast (PIM) traces to aid in PIM troubleshooting.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Start debug trace message output:


```
debug ip pim pimdbgtrace
```
3. Stop debug trace message output:


```
no debug ip pim pimdbgtrace
default debug ip pim pimdbgtrace
```
4. Configure the system to display trace messages forwarded by the device:


```
debug ip pim send-dbg-trace
```
5. Stop the system from displaying trace messages forwarded by the device:


```
no debug ip pim send-dbg-trace
```


- ```
default debug ip pim send-dbg-trace
```
6. Configure the system to display trace messages received by the device:
 

```
debug ip pim rcv-dbg-trace
```
  7. Stop the system from displaying trace messages received by the device:
 

```
no debug ip pim rcv-dbg-trace
```

```
default debug ip pim rcv-dbg-trace
```
  8. Configure the system to display hello messages forwarded or received by the device:
 

```
debug ip pim hello
```
  9. Stop the system from displaying hello messages forwarded or received by the device:
 

```
no debug ip pim hello
```

```
default debug ip pim hello
```
  10. Configure the system to display and log debug trace messages:
 

```
debug ip pim pimdbglog
```
  11. Stop the system from displaying and logging debug trace messages:
 

```
no debug ip pim pimdbglog
```

```
default debug ip pim pimdbglog
```
  12. Configure the system to display register messages forwarded or received by the device:
 

```
debug ip pim register
```
  13. Stop the system from displaying register messages forwarded or received by the device:
 

```
no debug ip pim register
```

```
default debug ip pim register
```
  14. Configure the system to display debug trace messages after an enabled message type, for example, hello or register, is received from a specific sender IP address:
 

```
debug ip pim source {A.B.C.D}
```

## Variable definitions

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

**Table 53: Variable definitions**

| Variable | Value                                                              |
|----------|--------------------------------------------------------------------|
| assert   | Displays the assert debug traces. The default is false (disabled). |
| bstrap   | Displays bootstrap debug traces. The default is false (disabled).  |

*Table continues...*

| Variable         | Value                                                                                       |
|------------------|---------------------------------------------------------------------------------------------|
| group {A.B.C.D}  | Displays debug traces from a specific group IP address. The default is 0.0.0.0 (disabled).  |
| hello            | Displays hello debug traces. The default is false (disabled).                               |
| joinprune        | Displays join and prune debug traces. The default is false (disabled).                      |
| pimdbglog        | Logs debug traces. The default is false (disabled).                                         |
| pimdbgtrace      | Displays PIM debug traces. The default is false (disabled).                                 |
| rcv-dbg-trace    | Displays trace messages received by the switch. The default is false (disabled).            |
| register         | If enabled, the system displays register debug traces. The default is false (disabled).     |
| regstop          | Displays register stop debug traces. The default is false (disabled).                       |
| rp-adv           | Displays RP advertisement debug traces. The default is false (disabled).                    |
| send-dbg-trace   | Displays trace messages forwarded by the switch. The default is false (disabled).           |
| source {A.B.C.D} | Displays debug traces from a specific source IP address. The default is 0.0.0.0 (disabled). |

## Determining the protocol configured on the added VLAN

Use this command to determine the protocol configured on the added VLAN.

The protocol configured on the added VLAN can be one of the following values:

- snoop
- snoop-spb
- route-spb
- pim

### Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Determine the protocol configured on the added VLAN:

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

The protocol displays under the Mode column of the command output.

## Example

Determine the protocol configured on the added VLAN:

```
Switch:1enable
Switch:1#show ip igmp interface

=====
IGMP Interface - GlobalRouter
=====
IF QUERY INTVL STATUS VERS. OPER VERS QUERIER QUERY WRONG JOINS ROBUST LASTMEM MODE
=====
V100 125 activ 2 2 0.0.0.0 100 0 0 2 10 routed-spb
=====
1 out of 1 entries displayed
```

## Variable definitions

Use the information in the following table to use the `show ip igmp interface` command.

| Variable                                                  | Value                                                                                                                                                                                                                      |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>gigabitethernet{slot/port[-slot/port][,...]}</code> | Specifies the slot/port.                                                                                                                                                                                                   |
| <code>vlan&lt;1-4084&gt;</code>                           | Specifies the VLAN ID in the range of 1 to 4084. VLAN IDs 1 to 4084 are configurable. The system reserves VLAN IDs 4085 to 4094 for internal use. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. |
| <code>vrfWORD&lt;1-16&gt;</code>                          | Specifies the VRF instance by the VRF name.                                                                                                                                                                                |
| <code>vrfidsWORD&lt;0-512&gt;</code>                      | Specifies the VRF ID for which to display statistics.                                                                                                                                                                      |

## Job aid

The following table shows the field descriptions for the `show ip igmp interface` command.

| Field       | Description                                                                                                                                                                                                                            |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IF          | Indicates the interfaces where IGMP is configured.                                                                                                                                                                                     |
| QUERY INTVL | Indicates the frequency at which the interface transmits IGMP host query packets.                                                                                                                                                      |
| STATUS      | Indicates the activation of a row that enables IGMP on the interface. The destruction of a row disables IGMP on the interface.                                                                                                         |
| VERS        | Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP. |
| OPER VERS   | Indicates the operational version of IGMP.                                                                                                                                                                                             |
| QUERIER     | Indicates the address of the IGMP querier on the IP subnet to which this interface is attached.                                                                                                                                        |

*Table continues...*

| Field          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| QUERY MAXRSPT  | Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.                                                                                                                                                                                                                                                                                                                                     |
| WRONG QUERY    | Indicates the number of queries received whose IGMP version does not match the IGMP Interface version. You must configure all routers on a LAN to run the same version of IGMP. Therefore, if the interface receives queries with the wrong version, a configuration error occurs.                                                                                                                                                                    |
| JOINS          | Indicates the number of times IGMP added a group membership on this interface.                                                                                                                                                                                                                                                                                                                                                                        |
| ROBUST         | Indicates the robustness variable, which you configure for the expected packet loss on a subnet. If packet loss is expected on a subnet, increase the robustness variable.                                                                                                                                                                                                                                                                            |
| LAST MEM QUERY | Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This does not apply if igmpInterface version is 1.                              |
| MODE           | Indicates the protocol configured on the VLAN added. <ul style="list-style-type: none"> <li>• snoop — Indicates IGMP snooping is enabled on a VLAN.</li> <li>• snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over SPBM for a Layer 2 VSN.)</li> <li>• routed-spb — Indicates IP multicast over SPBM is enabled on the Layer 3 VSN or for IP Shortcuts.</li> <li>• pim — Indicates PIM is enabled.</li> </ul> |

The following table shows the field descriptions for the `show ip igmp interface` command output if you use the optional parameters to specify a port, VLAN, or VRF.

**Table 54: show ip igmp interface command with optional parameters**

| Field               | Description                                                                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN ID or PORT NUM | Identifies the VLAN or port where IGMP is configured.                                                                                                                         |
| QUERY INTVL         | Indicates the frequency at which IGMP host query packets transmit on this interface.                                                                                          |
| QUERY MAX RESP      | Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.                                                             |
| ROBUST              | Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable. |

*Table continues...*

| Field                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VERSION                                     | Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.                                                                                                                                                                        |
| LAST MEMB QUERY                             | Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1. |
| PROXY SNOOP ENABLE                          | Indicates if proxy snoop is enabled on the interface.                                                                                                                                                                                                                                                                                                                                                         |
| SNOOP ENABLE                                | Indicates if snoop is enabled on the interface.                                                                                                                                                                                                                                                                                                                                                               |
| SSM SNOOP ENABLE                            | Indicates if SSM snoop is enabled on the interface.                                                                                                                                                                                                                                                                                                                                                           |
| FAST LEAVE ENABLE                           | Indicates if fast leave mode is enabled on the interface.                                                                                                                                                                                                                                                                                                                                                     |
| FAST LEAVE PORTS (VLAN parameter only)      | Indicates the set of ports that are enabled for fast leave.                                                                                                                                                                                                                                                                                                                                                   |
| VLAN ID or PORT NUM                         | Identifies the VLAN or port where IGMP is configured.                                                                                                                                                                                                                                                                                                                                                         |
| SNOOP QUERIER ENABLE (VLAN parameter only)  | Indicates if the IGMP Layer 2 Querier feature is enabled.                                                                                                                                                                                                                                                                                                                                                     |
| SNOOP QUERIER ADDRESS (VLAN parameter only) | Indicates the IP address of the IGMP Layer 2 Querier.                                                                                                                                                                                                                                                                                                                                                         |
| DYNAMIC DOWNGRADE VERSION                   | Indicates if the dynamic downgrade feature is enabled.                                                                                                                                                                                                                                                                                                                                                        |
| COMPATIBILITY MODE                          | Indicates if compatibility mode is enabled.                                                                                                                                                                                                                                                                                                                                                                   |
| EXPLICIT HOST TRACKING                      | Indicates if explicit host tracking is enabled to track all the source and group members.                                                                                                                                                                                                                                                                                                                     |

## Determining the data stream learned with IP multicast over SPBM on the VLAN

Use this command to determine the data stream learned when IP multicast over SPBM is configured on the VLAN.

### About this task

The following section shows sample output for the `show ip mroute route` command.

In this table, every stream uses one (\*,G) entry and x (S,G) entries, depending on how many servers forward traffic to the same group.

The 0.0.0.0 mask is always tied to a (\*,G) entry.

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

**Procedure**

1. Log on to the switch to enter User EXEC mode.
2. Determine the data stream learned:

```
show ip mroute route [vrf WORD <0-32>] [vrfids <0-255>]
```

**Example**

Determine the data stream learned:

```
Switch:1#show ip mroute route
=====
Mroute Route
=====
GROUP SOURCE SRCMASK UPSTREAM_NBR IF EXPIR PROT

233.252.0.119 0.0.0.0 0.0.0.0 203.0.113.207 v504 210 pimsm
233.252.0.119 192.0.2.165 255.255.255.0 203.0.113.207 v504 210 pimsm
233.252.0.119 198.51.100.165 255.255.255.0 198.51.100.204 v155 210 pimsm
233.252.0.119 198.51.100.166 255.255.255.0 198.51.100.204 v155 210 pimsm
233.252.0.113 0.0.0.0 0.0.0.0 203.0.113.208 v504 210 pimsm
233.252.0.113 198.51.100.165 255.255.255.0 203.0.113.208 v504 210 pimsm
```

**Job aid**

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

**Table 55: Job aid: show ip mroute route command output**

| Field        | Description                                                                                                                                                                                             |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GROUP        | Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface                                                                                                     |
| SOURCE       | Indicates the network address that, when combined with the corresponding value of ipMRouteNextHopSourceMask, identifies the sources for which this entry specifies a next hop on an outgoing interface. |
| SRCMASK      | Indicates the network mask, when combined with the corresponding value of ipMRouteNextHopSource, identifies the sources for which this entry specifies a next hop on an outgoing interface.             |
| UPSTREAM_NBR | Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is known.                         |
| IF           | Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject              |

*Table continues...*

| Field | Description                                                                                                                                                                                                                                                                                     |
|-------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|       | to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).                                                                                                                                                                                               |
| EXPIR | Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.                                                                                                                                                        |
| PROT  | Indicates the outgoing mechanism through which the switch learns this route. For IP multicast over SPBM, this value is sb-access or spb-network. Spb-access indicates the datastream learned was from the UNI ports. Spb-network indicates that the datastream learned was from the SPBM cloud. |

## Displaying the SPBM multicast database

You can determine the database used by the SPBM multicast module by using the following procedure.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Show the SPBM multicast database:

```
show isis spbm ip-multicast-route [all][detail][group {A.B.C.D}]
[vlan <0-4084>][vrf WORD<1-16>][vsid <1-16777215>]
```

### Example

Show the SPBM multicast database:

```
Switch(config)#show isis spbm ip-multicast-route
```

```
=====
 SPBM IP-MULTICAST FIB ENTRY INFO
=====
Source Group Data ISID BVLAN Source-BEB

192.2.0.1 233.252.0.246 16000001 101 EVP

Total Number of SPBM IP MULTICAST ROUTE Entries: 1
=====
```

## Variable definitions

Use the data in the following table to use the `show isis spbm ip-multicast-route` command.

| Variable | Value                                                  |
|----------|--------------------------------------------------------|
| all      | Displays all IP multicast over SPBM route information. |

*Table continues...*

| Variable                         | Value                                                                                                                                               |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| detail                           | Displays detailed IP multicast over SPBM route information.                                                                                         |
| group {A.B.C.D} source {A.B.C.D} | Displays information on the group IP address for the IP multicast over SPBM route. If you select source it will also display the source IP address. |
| vlan <0–4084>                    | Displays IP multicast over SPBM route information by VLAN.                                                                                          |
| vrf WORD<1–16>                   | Displays IP multicast over SPBM route information by VRF.                                                                                           |
| vsn-isid <1–16777215>            | Displays IP multicast over SPBM route information by I-SID.                                                                                         |

## Job aid

The following table describes fields for the `show isis spbm ip-multicast-route` command.

**Table 56: show isis spbm ip-multicast-route command**

| Field                                           | Description                                                                                          |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Source                                          | Specifies the IP address of the Global Routing Table.                                                |
| Group                                           | Specifies the IP multicast group for which this entry specifies a next hop on an outgoing interface. |
| Data ISID                                       | Specifies the VRF ID for the multicast route.                                                        |
| BVLAN                                           | Specifies the Backbone VLAN (B-VLAN).                                                                |
| Source-BEB                                      | Specifies the source Backbone Edge Bridge (BEB).                                                     |
| Total number of SPBM IP_MULTICAST Route entries | Specifies the number of SPBM IP multicast route entries.                                             |

## Troubleshooting IP multicast over SPBM for Layer 2 VSNs

If traffic is not moving properly, use the following checklist to determine the issue.

### Procedure

1. Ensure that all switch nodes in the network operate with the current release:
 

```
show software
```
2. If any ERS 8800 nodes exist in the network, ensure you upgrade them to the current release:
 

```
show software
```
3. Ensure that you create and enable SPBM infrastructure globally.
  - a. Ensure that SPBM is enabled globally:
 

```
show spbm
```



- b. Ensure that IS-IS is enabled globally:

```
show isis
```

- c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

```
show isis spbm
```

For more information on infrastructure and services configuration, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000*, NN46250-510.

- 4. Ensure that you enable the CFM configuration.

- a. Ensure a CFM maintenance-association exists:

```
show cfm maintenance-association
```

- b. Ensure a CFM maintenance-domain exists:

```
show cfm maintenance-domain
```

- c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

```
show cfm maintenance-endpoint
```

- 5. Ensure a Customer VLAN (C-VLAN) exists and ensure you add UNI ports to the C-VLAN.

- a. Display C-VLAN information:

```
show vlan i-sid
```

- b. Display ports for the C-VLAN:

```
show vlan members port {slot/port [-slot/port][,...]}
```

- c. Display NNI and UNI receivers:

```
show isis spbm ip-multicast-route detail
```

- 6. Ensure that you assign the same I-SID to the C-VLAN on all of the BEBs where you configure the C-VLAN:

```
show vlan i-sid
```

- 7. Ensure that you enable IP multicast over SPBM globally:

```
show isis spbm
```

- 8. Enter Privileged EXEC mode:

```
enable
```

- 9. Ensure the you enable IGMP Snooping on the C-VLAN on all of the Backbone Edge Bridges (BEBs). Ensure the protocol configured on the VLAN added is snoop-spb in the MODE column, which indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over SPBM for a Layer 2 VSN):

```
show ip igmp interface
```

10. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

```
show ip igmp snoop-trace
show ip igmp interface
```

11. Enter Global Configuration mode:

```
enable
configure terminal
```

12. Configure an IGMP Querier address on the C-VLAN if the access Layer 2 switch does not recognize a 0.0.0.0 IP Querier address:

```
ip igmp snoop-querier-addr <A.B.C.D>
```

13. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

```
show ip igmp interface
```

---

## Troubleshooting IP multicast over SPBM for Layer 3 VSNs

If traffic is not moving properly, use the following checklist to determine the issue.

### Procedure

1. Ensure that all switch nodes in the network operate with the current release:

```
show software
```

2. If ERS 8800 nodes exist in the network, ensure you upgrade them to the current release:

```
show software
```

3. Ensure that you create and enable SPBM infrastructure globally.

- a. Ensure that SPBM is enabled globally:

```
show spbm
```

- b. Ensure that IS-IS is enabled globally:

```
show isis
```

- c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

```
show isis spbm
```

For more information on infrastructure and services configuration, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000*, NN46250-510.

4. Ensure that you enable the CFM configuration.

- a. Ensure a CFM maintenance-association exists:

```
show cfm maintenance-association
```

- b. Ensure a CFM maintenance-domain exists:

```
show cfm maintenance-domain
```

- c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

```
show cfm maintenance-endpoint
```

5. Ensure the following on all the Backbone Edge Bridges (BEBs) where the Layer 3 VSN is present.

- a. Ensure that you enable IP multicast globally:

```
show isis spbm
```

- b. Ensure that you create an IPVPN for the VRF:

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

- c. Ensure that you assign an I-SID to the VRF:

```
show isis spbm ip-multicast-route all
```

- d. Ensure that you enable the MVPN:

```
show ip vrf mvpn
```

6. On the VLANs that need Layer 3 VSN IP multicast over SPBM routing, create an IP interface on the VLAN if one does not exist. The address should be on the same subnet as the IGMP hosts connected to the VLAN. Also, ensure that you enable IP multicast over SPBM.

7. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4084>
```

8. Create an IP interface on the VLAN and enable IP multicast over SPBM:

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

```
ip spb-multicast enable
```

9. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

```
show ip igmp snoop-trace
```

```
show ip igmp interface
```

10. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

```
show ip igmp interface
```

---

## Troubleshooting IP multicast over SPBM for IP Shortcuts

If traffic is not moving properly, use the following checklist to determine the issue.

### Procedure

1. Ensure that all switch nodes in the network operate with the current release:

```
show software
```

2. Ensure that all ERS 8800 nodes in the network have the current release:

```
show software
```

3. Ensure that you create and enable SPBM infrastructure globally.

- a. Ensure that SPBM is enabled globally:

```
show spbm
```

- b. Ensure that IS-IS is enabled globally:

```
show isis
```

- c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

```
show isis spbm
```

For more information on infrastructure and services configuration, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000*, NN46250-510.

4. Ensure that you enable the CFM configuration.

- a. Ensure a CFM maintenance-association exists:

```
show cfm maintenance-association
```

- b. Ensure a CFM maintenance-domain exists:

```
show cfm maintenance-domain
```

- c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

```
show cfm maintenance-endpoint
```

5. Ensure the following on all BEBs where you want IP multicast over SPBM. Ensure that you enable IP multicast over SPBM globally:

```
show isis spbm
```

6. On the VLANs that need Layer 3 VSN IP multicast over SPBM routing, create an IP interface on the VLAN if one does not exist. The address should be on the same subnet as the IGMP hosts connected to the VLAN. Also, ensure that you enable IP multicast over SPBM. Create an IP interface on the VLAN and enable IP multicast over SPBM.

7. Enter VLAN Interface Configuration mode:

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

8. Create an IP interface on the VLAN and enable IP multicast over SPBM:

```
ip address <A.B.C.D>
ip spb-multicast enable
```

9. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

```
show ip igmp snoop-trace
show ip igmp interface
```

10. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

```
show ip igmp interface
```

---

## Defining the IS-IS trace flag for IP multicast

Define the IS-IS trace flag for IP multicast.

### Procedure

1. Log on to the switch to enter User EXEC mode.
2. Define the IS-IS trace flag for IP multicast:

```
trace flags isis [set ip-multicast] [remove ip-multicast]
```

---

## Multicast routing troubleshooting using EDM

Use the information in this section to help you troubleshoot multicast routing problems using Enterprise Device Manager (EDM).

---

### Viewing IGMP interface information

Use the Interface tab to view the IGMP interface table. You can use this command to determine the protocol configured on the added VLAN.

The protocol configured on the added VLAN can be one of the following values:

- snoop

- snoop-spb
- route-spb
- pim

**About this task**


If an interface does not use an IP address, it does not appear in the IGMP table. If an interface uses an IP address, but neither IGMP snoop or PIM is enabled, the interface appears as inactive in the Status field.

**Procedure**

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

**Interface field descriptions**

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

| Name                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IfIndex</b>              | Shows the interface where IGMP is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>QueryInterval</b>        | Configures the frequency (in seconds) at which the IGMP host query packets transmit on the interface. The range is from 1–65535 and the default is 125.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Status</b>               | Shows the IGMP row status. If an interface uses an IP address and PIM-SM is enabled, the status is active. Otherwise, it is notInService.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Version</b>              | Configures the version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.                                                                                                                                                                                                                                                                                                         |
| <b>OperVersion</b>          | Shows the version of IGMP that currently runs on this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Querier</b>              | Shows the address of the IGMP querier on the IP subnet to which this interface attaches.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>QueryMaxResponseTime</b> | <p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1.</p> <p>Smaller values allow a router to prune groups faster. The range is from 0–255, and the default is 100 tenths of a second (equal to 10 seconds.)</p> <p> <b>Important:</b><br/>You must configure this value lower than the QueryInterval.</p> |
| <b>WrongVersionQueries</b>  | Shows the number of queries received with an IGMP version that does not match the interface. You must configure all routers on a LAN to run the same version of IGMP. If the interface receives                                                                                                                                                                                                                                                                                                             |

*Table continues...*

| Name                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                    | queries with the wrong version, this value indicates a version mismatch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Joins</b>                       | Shows the number of times this interface added a group membership, which is the same as the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Robustness</b>                  | <p>Tunes for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, increase the robustness value.</p> <p>The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.</p>                                                                                                                                                                                                                                           |
| <b>LastMembQueryIntvl</b>          | <p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of second. Avaya recommends that you configure this parameter to values greater than 3. If you do not need a fast leave process, Avaya recommends values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)</p> |
| <b>OtherQuerierPresent Timeout</b> | Shows the length of time that must pass before a multicast router determines that no other querier exists. If the local router is the querier, the value is 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>FlushAction</b>                 | <p>Configures the flush action to one of the following:</p> <ul style="list-style-type: none"> <li>• none</li> <li>• flushGrpMem</li> <li>• flushMrouter</li> <li>• flushSender</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>RouterAlertEnable</b>           | <p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p><b>!</b> <b>Important:</b></p> <p>To maximize network performance, Avaya recommends that you configure this parameter according to the version of IGMP currently in use.</p> <ul style="list-style-type: none"> <li>• IGMPv1—Disable</li> <li>• IGMPv2—Enable</li> </ul>                                                                                                                                     |

*Table continues...*

| Name                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                   | <ul style="list-style-type: none"> <li>IGMPv3—Enable</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>SsmSnoopEnable</b>             | Enables SSM snoop.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>SnoopQuerierEnable</b>         | Enables IGMP Layer 2 Querier.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>SnoopQuerierAddr</b>           | Enables the IGMP Layer 2 Querier address.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>ExplicitHostTrackingEnable</b> | Enables or disables IGMPv3 to track hosts for each channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.                                                                                                                                                                                                                                                                             |
| <b>McastMode</b>                  | <p>Indicates the protocol configured on the VLAN.</p> <ul style="list-style-type: none"> <li>snoop — Indicates IGMP snooping is enabled on a VLAN.</li> <li>snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over SPBM for a Layer 2 VSN.)</li> <li>routed-spb — Indicates IP multicast over SPBM is enabled on the Layer 3 VSN or for IP Shortcuts.</li> <li>pim — Indicates PIM is enabled.</li> </ul> |

## Viewing IGMP snoop trace information

### About this task

View the multicast group trace to track the data flow path of multicast streams.

### Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.
2. Click **IGMP**.
3. Click the **Snoop Trace** tab.

## Snoop Trace field descriptions

Use the data in the following table to use the **Snoop Trace** tab.

| Name           | Description                                                           |
|----------------|-----------------------------------------------------------------------|
| <b>GrpAddr</b> | Displays the IP multicast address of the group traversing the router. |
| <b>SrcAddr</b> | Displays the IP source address of the multicast group.                |
| <b>OutVlan</b> | Displays the egress VLAN ID for the multicast group.                  |
| <b>InPort</b>  | Displays the ingress port for the multicast group.                    |
| <b>InVlan</b>  | Displays the ingress VLAN ID for the multicast group.                 |

*Table continues...*



| Name    | Description                                                 |
|---------|-------------------------------------------------------------|
| OutPort | Displays the egress port of the multicast group.            |
| Type    | Displays the port type on which the snoop entry is learned. |

## Viewing IGMP group information

View information about IGMP groups to see the current group operation on the switch.

### About this task

#### \* Note:

The following procedure displays the dynamically learned IGMP groups. **IP > IGMP > Static** displays statically configured IGMP groups. This is in contrast to the ACLI command **show ip igmp group**, which displays both dynamically learned and statically configured IGMP groups, and the ACLI command **show ip igmp static**, which displays only the statically configured groups.

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

### Procedure

1. In the navigation pane, expand the following folders: **Configuration > IP**.
2. Click **IGMP**.
3. Click the **Groups** tab.

## Groups field descriptions

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

| Name       | Description                                                                                                                        |
|------------|------------------------------------------------------------------------------------------------------------------------------------|
| IpAddress  | Shows the multicast group address (Class D). A group address can be the same for many incoming ports.                              |
| Members    | Shows the IP address of the host that issues the membership report to this group.                                                  |
| InPort     | Shows the port that receives the group membership report.                                                                          |
| IfIndex    | Shows a unique value that identifies a physical interface or a logical interface (VLAN) that receives the membership report.       |
| Expiration | Shows the time left before the group report expires on this port. This variable is updated after the port receives a group report. |

---

## Viewing multicast group sources

With the Sources tab, you can view all the sources on the subnet that send to the particular group selected in the Mroute-HW table.

### Procedure

1. In the navigation tree, expand the following folders: **IP > Multicast**.
2. Click the **Mroute-HW** tab.
3. Click any row in the table.
4. Click **Sources**.

## Sources field descriptions

Use the information in the following table to help you understand the **Source** tab fields.

| Name                 | Description                                                                                                                                     |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SourceAddress</b> | The IP addresses of the sources on this particular subnet sending traffic to the multicast group for the selected entry in the Mroute-HW table. |
| <b>IngressPort</b>   | The corresponding ingress port in the multicast stream selected from the Mroute-HW table.                                                       |

---

## Viewing multicast routes by egress VLAN

With the Egress VLANs tab, you can view the egress VLANs for the streams corresponding to the selected entry in the Mroute-Hw table.

### Procedure

1. In the navigation tree, expand the following folders: **IP > Multicast**.
2. Click the **Mroute-HW** tab.
3. Click any row in the table.
4. Click **EgressVlans**.

## EgressVlans field descriptions

Use the information in the following table to help you use the **EgressVlans** tab.

| Name                   | Description                                                                                    |
|------------------------|------------------------------------------------------------------------------------------------|
| <b>EgressVlan</b>      | All the egress VLANs for the particular multicast stream selected from the Mroute-HW table.    |
| <b>EgressVlanPorts</b> | The corresponding ports for the particular multicast stream selected from the Mroute-HW table. |

## Determining the data stream learned when IP multicast over SPBM is configured on the VLAN

Use the following procedure to determine the data stream learned when IP multicast is configured on the VLAN.

### Procedure

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

### Multicast field descriptions

Use the information in the following table to help you use the **Multicast** tab.

| Field                   | Description                                                                                                                                                                                                                                                                                      |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group</b>            | Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.                                                                                                                                                                                             |
| <b>Source</b>           | Indicates the network address that, when combined with the corresponding value of ipMRouteNextHopSourceMask, identifies the sources for which this entry specifies a next hop on an outgoing interface.                                                                                          |
| <b>SourceMask</b>       | Indicates the network mask, when combined with the corresponding value of ipMRouteNextHopSource, identifies the sources for which this entry specifies a next hop on an outgoing interface.                                                                                                      |
| <b>UpstreamNeighbor</b> | Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is known.                                                                                                                  |
| <b>Interface</b>        | Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).     |
| <b>ExpiryTime</b>       | Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.                                                                                                                                                         |
| <b>Protocol</b>         | Indicates the outgoing mechanism through which the switch learns this route. For IP multicast over SPBM, this value is spb-access or spb-network. Spb-access indicates the datastream learned was from the UNI ports. Spb-network indicates that the datastream learned was from the SPBM cloud. |

## Showing the SPBM multicast database

Determine the database used by the SPBM multicast module.

### Procedure

1. From the navigation tree, expand the following folders: **Configuration > ISIS > SPBM**.
2. Click the **IpMcastRoutes** tab.

## IpMcastRoutes field descriptions

Use the information in the following table to use the **IpMcastRoutes** tab.

| Name                 | Description                                                                 |
|----------------------|-----------------------------------------------------------------------------|
| <b>Vsnlsid</b>       | Specifies the VSN I-SID.                                                    |
| <b>Group</b>         | Specifies the group IP address for the IP multicast route.                  |
| <b>Source</b>        | Specifies the IP address where the IP multicast route originated from.      |
| <b>SourceBeb</b>     | Specifies the Source Backbone Edge Bridge (BEB) for the IP multicast route. |
| <b>VlanId</b>        | Specifies the VLAN ID.                                                      |
| <b>VrfName</b>       | Specifies the VRF name.                                                     |
| <b>Datalsid</b>      | Specifies the VRF ID for the multicast route.                               |
| <b>Type</b>          | Specifies the type for the IP multicast route.                              |
| <b>Bvlan</b>         | Specifies the Backbone VLAN (B-VLAN).                                       |
| <b>NniInterfaces</b> | Specifies the Network-to-Network Interface ports.                           |

# Chapter 14: Safety Messages

This section describes the different precautionary notices used in the Avaya Virtual Services Platform 9000 documentation. This section also contains precautionary notices that you must read for safe operation of Avaya Virtual Services Platform 9000.

## Notices

Notice paragraphs alert you about issues that require your attention. The following sections describe the types of notices. For a list of safety messages used in a document and their translations, see the Translations of safety messages chapter.

### Attention Notice

#### Important:

An attention notice provides important information regarding the installation and operation of Avaya products.

### Caution ESD Notice

#### Electrostatic alert:

##### **ESD**

ESD notices provide information about how to avoid discharge of static electricity and subsequent damage to Avaya products.

#### Electrostatic alert:

##### **ESD (décharge électrostatique)**

La mention ESD fournit des informations sur les moyens de prévenir une décharge électrostatique et d'éviter d'endommager les produits Avaya.

#### Electrostatic alert:

##### **ACHTUNG ESD**

ESD-Hinweise bieten Information dazu, wie man die Entladung von statischer Elektrizität und Folgeschäden an Avaya-Produkten verhindert.

#### Electrostatic alert:

##### **PRECAUCIÓN ESD (Descarga electrostática)**

El aviso de ESD brinda información acerca de cómo evitar una descarga de electricidad estática y el daño posterior a los productos Avaya.

 **Electrostatic alert:**

**CUIDADO ESD**

Os avisos do ESD oferecem informações sobre como evitar descarga de eletricidade estática e os conseqüentes danos aos produtos da Avaya.

 **Electrostatic alert:**

**ATTENZIONE ESD**

Le indicazioni ESD forniscono informazioni per evitare scariche di elettricità statica e i danni correlati per i prodotti Avaya.

**Caution Notice**

 **Caution:**

Caution notices provide information about how to avoid possible service disruption, loss of data, or harm to software.

 **Caution:**

**PELIGRO**

Los avisos de peligro proporcionan información sobre cómo evitar una posible interrupción del servicio, pérdida de datos o daño al software.

 **Caution:**

**ACHTUNG**

In diesen Hinweisen erfahren Sie, wie Sie Dienstunterbrechungen, Datenverlust oder Beeinträchtigungen der Software vermeiden können.

 **Caution:**

**MISE EN GARDE**

Les avis de mise en garde fournissent des informations indiquant comment éviter tout risque d'interruption de service, de perte de données ou de détérioration du logiciel.

 **Caution:**

**CUIDADO**

Avisos de cuidado fornecem informações sobre como evitar possíveis interrupções de serviço, perda de dados ou danos ao software.

 **Caution:**

**ATTENZIONE**

Un avvertimento di attenzione fornisce le informazioni su come evitare situazioni che potrebbero causare danni al software, perdita di dati o interruzione del servizio.

**Warning Notice**

 **Warning:**

Warning notices provide information about how to avoid harm to hardware or equipment.

 **Warning:**

**AVERTISSEMENT**

Les avis d'avertissements fournissent des informations indiquant comment éviter de détériorer le matériel ou un équipement.

 **Warning:**

**WARNUNG**

In Warnhinweisen erfahren Sie, wie Sie Beschädigungen der Hardware oder anderer Geräte vermeiden können.

 **Warning:**

**ADVERTENCIA**

Los avisos de advertencia proporcionan información sobre cómo evitar daño al hardware o al equipo.

 **Warning:**

**ADVERTÊNCIA**

Avisos de advertência fornecem informações sobre como evitar danos aos equipamentos.

 **Warning:**

**Avvertenza**

Un'avvertenza richiama fornisce le informazioni su come evitare situazione che potrebbero danneggiare l'hardware o l'apparecchiatura.

**Danger High Voltage Notice**

 **Voltage:**

Danger—High Voltage notices provide information about how to avoid a situation or condition that can cause serious personal injury or death from high voltage or electric shock.

 **Voltage:**

La mention Danger—Tension élevée fournit des informations sur les moyens de prévenir une situation ou une condition qui pourrait entraîner un risque de blessure grave ou mortelle à la suite d'une tension élevée ou d'un choc électrique.

 **Voltage:**

**GEFAHR**

Hinweise mit „Vorsicht – Hochspannung“ bieten Informationen dazu, wie man Situationen oder Umstände verhindert, die zu schweren Personenschäden oder Tod durch Hochspannung oder Stromschlag führen können.

 **Voltage:**

**PELIGRO**

Los avisos de Peligro-Alto voltaje brindan información acerca de cómo evitar una situación o condición que cause graves lesiones a personas o la muerte, a causa de una electrocución o de una descarga de alto voltaje.

 **Voltage:**

**PERIGO**

Avisos de Perigo—Alta Tensão oferecem informações sobre como evitar uma situação ou condição que possa causar graves ferimentos ou morte devido a alta tensão ou choques elétricos.

 **Voltage:**

**PERICOLO**

Le indicazioni Pericolo—Alta tensione forniscono informazioni per evitare situazioni o condizioni che potrebbero causare gravi danni alle persone o il decesso a causa dell'alta tensione o di scosse elettriche.

**Danger Notice**

 **Danger:**

Danger notices provide information about how to avoid a situation or condition that can cause serious personal injury or death.

 **Danger:**

La mention Danger fournit des informations sur les moyens de prévenir une situation ou une condition qui pourrait entraîner un risque de blessure grave ou mortelle.

 **Danger:**

**GEFAHR**

Gefahrenhinweise stellen Informationen darüber bereit, wie man Situationen oder Umständen verhindert, die zu schweren Personenschäden oder Tod führen können.

 **Danger:**

**PELIGRO**

Los avisos de Peligro brindan información acerca de cómo evitar una situación o condición que pueda causar lesiones personales graves o la muerte.

 **Danger:**

**PERIGO**

Avisos de perigo oferecem informações sobre como evitar uma situação ou condição que possa causar graves ferimentos ou morte.

 **Danger:**

**PERICOLO**

Le indicazioni di pericolo forniscono informazioni per evitare situazioni o condizioni che potrebbero causare gravi danni alle persone o il decesso.



# Glossary

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>attenuation</b>                             | The decrease in signal strength in an optical fiber caused by absorption and scattering.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Avaya command line interface (ACLI)</b>     | A textual user interface. When you use ACLI, you respond to a prompt by typing a command. After you enter the command, you receive a system response.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Backbone Core Bridge (BCB)</b>              | Backbone Core Bridges (BCBs) form the core of the SPBM network. The BCBs are SPBM nodes that do not terminate the VSN services. BCBs forward encapsulated VSN traffic based on the Backbone MAC Destination Address (B-MAC-DA). A BCB can access information to send that traffic to any Backbone Edge Bridges (BEBs) in the SPBM backbone.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Backbone Edge Bridge (BEB)</b>              | Backbone Edge Bridges (BEBs) are SPBM nodes where Virtual Services Networks (VSNs) terminate. BEBs handle the boundary between the core MAC-in-MAC Shortest Path Bridging MAC (SPBM) domain and the edge customer 802.1Q domain. A BEB node performs 802.1ah MAC-in-MAC encapsulation and decapsulation for the Virtual Services Network (VSN).                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Backbone MAC (B-MAC)</b>                    | Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation encapsulates customer MAC addresses in Backbone MAC (B-MAC) addresses. MAC-in-MAC encapsulation defines a B-MAC-DA and B-MAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that SPBM uses for delivery from end to end. As the MAC header stays the same across the network, no need exists to swap a label or perform a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end. In Shortest Path Bridging MAC (SPBM), each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network. |
| <b>Backbone VLAN identifier (B-VID)</b>        | The Backbone VLAN identifier (B-VID) indicates the Shortest Path Bridging MAC (SPBM) B-VLAN associated with the SPBM instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Complete Sequence Number Packets (CSNP)</b> | Complete Sequence Number Packets (CSNP) contain the most recent sequence numbers of all Link State Packets (LSPs) in the database. When all routers update their LSP database, synchronization is complete.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Connectivity Fault Management (CFM)</b> | Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) network. CFM operates at Layer 2 and provides the equivalent of ping and traceroute. IEEE 802.1ag Connectivity Fault Management (CFM) divides or separates a network into administrative domains called Maintenance Domains (MD).                                                                                                                                                                                                                                                                                             |
| <b>Control Processor (CP) module</b>       | The Control Processor module runs all high level protocols (BGP, OSPF) and distributes the results (routing updates) to the rest of the system. The CP manages and configures the IO and Switch Fabric modules, and maintains and monitors the health of the chassis.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Customer MAC (C-MAC)</b>                | For customer MAC (C-MAC) addresses, which is customer traffic, to forward across the service provider back, SPBM uses IEEE 802.1ah Provider Backbone Bridging MAC-in-MAC encapsulation. The system encapsulates C-MAC addresses within a backbone MAC (B-MAC) address pair made up of a B-MAC destination address (B-MAC-DA) and a B-MAC source address (B-MAC-SA).                                                                                                                                                                                                                                                                                                        |
| <b>cyclic redundancy check (CRC)</b>       | Ensures frame integrity is maintained during transmission. The CRC performs a computation on frame contents before transmission and on the receiving device. The system discards frames that do not pass the CRC.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>designated router (DR)</b>              | A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group. |
| <b>Dynamic Random Access Memory (DRAM)</b> | A read-write random-access memory, in which the digital information is represented by charges stored on the capacitors and must be repeatedly replenished to retain the information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Electrostatic Discharge (ESD)</b>       | The discharge of stored static electricity that can damage electronic equipment and impair electrical circuitry that results in complete or intermittent failures.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Enterprise Device Manager (EDM)</b>     | A Web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.                                                                                                                                                                                                                                                                                                                                                                                                                                    |

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>forwarding database (FDB)</b>                          | A database that maps a port for every MAC address. If a packet is sent to a specific MAC address, the switch refers to the forwarding database for the corresponding port number and sends the data packet through that port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Generalized Regular Expression Parser (grep)</b>       | A Unix command used to search files for lines that match a certain regular expression (RE).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>I/O cooling module (9012FC)</b>                        | The I/O cooling module is a hot swappable fan tray used to cool the I/O and CP modules in the Virtual Services Platform 9012.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>I/O module</b>                                         | An I/O module is a module that provides network connectivity for various media (sometimes called Layer 0) and protocol types. I/O modules are also called Ethernet modules.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Intermediate System to Intermediate System (IS-IS)</b> | <p>Intermediate System to Intermediate System (IS-IS) is a link-state, interior gateway protocol. ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System to Intermediate System (IS-IS). IS-IS operation is similar to Open Shortest Path First (OSPF).</p> <p>In Shortest Path Bridging MAC (SPBM) networks, IS-IS discovers network topology and builds shortest path trees between network nodes that IS-IS uses for forwarding unicast traffic and determining the forwarding table for multicast traffic. SPBM employs IS-IS as the interior gateway protocol and implements additional Type-Length-Values (TLVs) to support additional functionality.</p> |
| <b>Internet Assigned Numbers Authority (IANA)</b>         | The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Internet Control Message Protocol (ICMP)</b>           | A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Internet Group Management Protocol (IGMP)</b>          | IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Internet Protocol multicast (IPMC)</b>                 | The technology foundation for audio and video streaming, push applications, software distribution, multipoint conferencing, and proxy and caching solutions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>interswitch trunking (IST)</b>                         | A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.                                                                                                                                                                                                                                                                                                                                                                                              |

|                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IS-IS Hello packets</b>                    | Intermediate System to Intermediate System (IS-IS) uses Hello packets to initialize and maintain adjacencies between neighboring routers. IS-IS Hello packets contain the IP address of the interface over which the Hello transmits. These packets are broadcast to discover the identities of neighboring IS-IS systems and to determine whether the neighbor is a Level 1 router.                                                                                                 |
| <b>Layer 1</b>                                | Layer 1 is the Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interacts with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.                                                                                                                                                                                                                                                                  |
| <b>Layer 2</b>                                | Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Layer 2 Virtual Services Network</b>       | The Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).                                                                                                                         |
| <b>Layer 3</b>                                | Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Layer 3 Virtual Services Network</b>       | The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN). |
| <b>Layer 4</b>                                | The Transport Layer of the OSI model. An example of a Layer 4 protocol is Transmission Control Protocol (TCP).                                                                                                                                                                                                                                                                                                                                                                       |
| <b>light emitting diode (LED)</b>             | A semiconductor diode that emits light when a current passes through it.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Link State Packets (LSP)</b>               | Link State Packets (LSP) contain information about the state of adjacencies or defined and distributed static routes. Intermediate System to Intermediate System (IS-IS) exchanges this information with neighboring IS-IS routers at periodic intervals. Every router in the domain has an identical link state database and each runs shortest path first to calculate routes.                                                                                                     |
| <b>Link State Protocol Data Unit (LSPDUs)</b> | Link State Protocol Data Unit is similar to a Link State Advertisement in Open Shortest Path First (OSPF). Intermediate System to Intermediate System (IS-IS) runs on all nodes of Shortest Path Bridging-MAC (SPBM).                                                                                                                                                                                                                                                                |

Since IS-IS is the basis of SPBM, the device must first form the IS-IS adjacency by first sending out hellos and then Link State Protocol Data Units. After the hellos are confirmed both nodes send Link State Protocol Data Units (LSPDUs) that contain connectivity information for the SPBM node. These nodes also send copies of all other LSPDUs they have in their databases. This establishes a network of connectivity providing the necessary information for each node to find the best and proper path to all destinations in the network.

### **link trace message**

The link trace message (LTM) is often compared to traceroute. A MEP transmits the LTM packet. This packet specifies the target MAC address of an MP, which is the SPBM system id or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR. LTM contains:

- Time to live (TTL)
- Transaction Identifier
- Originator MAC address
- Target MAC address

### **link-state database (LSDB)**

A database built by each OSPF router to store LSA information. The router uses the LSDB to calculate the shortest path to each destination in the autonomous system (AS), with itself at the root of each path.

### **Local Area Network (LAN)**

A data communications system that lies within a limited spatial area, uses a specific user group and topology, and can connect to a public switched telecommunications network (but is not one).

### **Loopback Messages (LBM)**

A Loopback Message (LBM) is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a Maintenance End Point (MEP) or Maintenance Intermediate Point (MIP), but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR. The system achieves fault verification through the use of Loopback Messages (LBM).

### **Loopback Response (LBR)**

Loopback Response (LBR) is the response from a Maintenance Point (MP).

### **MAC-in-MAC encapsulation**

MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that the device uses for delivery from end to end. As the MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end.

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Maintenance Associations (MA)</b>     | Maintenance Associations (MA) are administrative associations in a network that is divided by the 802.1ag Connectivity Fault Management (CFM) feature. CFM groups MAs within Maintenance Domains. Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network. |
| <b>Maintenance Domains (MD)</b>          | Maintenance Domains (MD) are administrative domains that divides a network by the 802.1ag Connectivity Fault Management (CFM) feature. Each MD is further subdivided into logical groupings called Maintenance Associations (MA). Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.                                                                                                             |
| <b>Maintenance Points (MP)</b>           | Maintenance Points (MP) are a demarcation point on an interface that participates in Connectivity Fault Management (CFM) within a Maintenance Domain (MD). There are two types of MP: Maintenance End Point (MEP) and Maintenance Intermediate Point (MIP). Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging MAC (SPBM) Network.                                                                                   |
| <b>management information base (MIB)</b> | The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>mask</b>                              | A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.                                                                                                                                                                                                                                                                                                                                             |
| <b>maximum transmission unit (MTU)</b>   | The largest number of bytes in a packet—the maximum transmission unit of the port.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Media Access Control (MAC)</b>        | MAC arbitrates access to and from a shared medium.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>mirrored port</b>                     | The port to mirror. The port is also called the source port.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>mirroring multilink trunk</b>         | The multilink trunk to which the system mirrors the traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>mirroring port</b>                    | The port to which the system mirrors all traffic, also referred to as the destination port.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>mirroring VLAN</b>                    | The virtual Local Area Network (VLAN) to which the system mirrors the traffic.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>multicast group ID (MGID)</b>         | The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a                                                                                                                                                                                                                                                                                                                                                |

specific port number, the switch directs the data to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs.

|                                                      |                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MultiLink Trunking (MLT)</b>                      | A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.                                                                                                      |
| <b>next hop</b>                                      | The next hop to which a packet can be sent to advance the packet to the destination.                                                                                                                                                                                                                                                                           |
| <b>nonbroadcast multiaccess (NBMA)</b>               | Interconnects multiple devices over a broadcast network through point-to-point links. NBMA reduces the number of IP addresses required for point-to-point connections.                                                                                                                                                                                         |
| <b>Open Systems Interconnection (OSI)</b>            | A suite of communication protocols, network architectures, and network management standards produced by the International Organization for Standardization (ISO). OSI-compliant systems can communicate with other OSI-compliant systems for a meaningful exchange of information.                                                                             |
| <b>Packet Capture Tool (PCAP)</b>                    | A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.                                                                                                                                                                               |
| <b>Partial Sequence Number Packets (PSNP)</b>        | Partial Sequence Number Packets (PSNP) are requests for missing Link State Packets (LSPs). When a receiving router detects a missing LSP, it sends a PSNP to the router that sent the Complete Sequence Number Packets (CSNP).                                                                                                                                 |
| <b>port mirroring</b>                                | A feature that sends received or transmitted traffic to a second destination.                                                                                                                                                                                                                                                                                  |
| <b>Protocol Data Units (PDUs)</b>                    | A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.                                                                                                                                                                               |
| <b>Provider Backbone Bridge (PBB)</b>                | To forward customer traffic across the service-provider backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a B-MAC-DA and B-MAC-SA to identify the backbone source and destination addresses. |
| <b>quad small form-factor pluggable plus (QSFP+)</b> | The quad small form-factor pluggable plus (QSFP+) is a transceiver used for data communications applications, which connects a module to a fiber optic cable.                                                                                                                                                                                                  |
| <b>quality of service (QoS)</b>                      | QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can                                                                                                                                                                                                                   |

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                            | configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>remote mirror source (RMS)</b>          | The port that generates the mirrored encapsulated traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>remote mirror target (RMT)</b>          | The port that decapsulates the remote mirror traffic and transmits it out of the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>remote mirroring</b>                    | A mirroring port that encapsulates traffic into a Layer 2 header and transmits it to a remote mirror target (RMT) for decapsulation. The packet transmits over a Layer 2 network and preserves the original packet.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>remote monitoring (RMON)</b>            | A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.                                                                                                                                                                                                                                                                                                                            |
| <b>route table manager (RTM)</b>           | Determines the best route to a destination based on reachability, route preference, and cost.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Routing Information Protocol (RIP)</b>  | A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.                                                                                                                                                                                                                                                                                                                        |
| <b>Secure Shell (SSH)</b>                  | SSH uses encryption to provide security for remote logons and data transfer over the Internet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Secure Sockets Layer (SSL)</b>          | An Internet security encryption and authentication protocol for secure point-to-point connections over the Internet and intranets, especially between clients and servers.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Service Instance Identifier (I-SID)</b> | The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 24 bits. SPBM uses this I-SID to identify and transmit any virtualized traffic in an encapsulated SPBM frame. SPBM uses I-SIDs to virtualize VLANs (Layer 2 Virtual Services Network [VSN]) or VRFs (Layer 3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone. With Layer 2 VSNs, you associate the I-SID with a customer VLAN, which is then virtualized across the backbone. With Layer 3 VSNs, you associate the I-SID with a customer VRF, which is also virtualized across the backbone. |
| <b>Shortest Path Bridging (SPB)</b>        | Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses                                                                                                                                                                                                                             |



the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.

**Shortest Path Bridging MAC (SPBM)**

Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.

**Simple Loop Prevention Protocol (SLPP)**

Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).

**Simple Network Management Protocol (SNMP)**

SNMP administratively monitors network performance through agents and management stations.

**small form-factor pluggable (SFP)**

A hot-swappable input and output enhancement component used with Avaya products to allow gigabit Ethernet ports to link with other gigabit Ethernet ports over various media types.

**small form-factor pluggable plus (SFP+)**

SFP+ transceivers are similar to SFPs in physical appearance but SFP+ transceivers provide Ethernet at 10 gigabits per second (Gbps).

**spanning tree**

A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.

**Spanning Tree Group (STG)**

A collection of ports in one spanning-tree instance.

**Split MultiLink Trunking (SMLT)**

An Avaya extension to IEEE 802.1AX (link aggregation), provides nodal and link failure protection and flexible bandwidth scaling to improve on the level of Layer 2 resiliency.

**time-to-live (TTL)**

The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.

|                                                          |                                                                                                                                                                                                                                                                                |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Trivial File Transfer Protocol (TFTP)</b>             | A protocol that governs transferring files between nodes without protection against packet loss.                                                                                                                                                                               |
| <b>trunk</b>                                             | A logical group of ports that behaves like a single large port.                                                                                                                                                                                                                |
| <b>unshielded twisted pair (UTP)</b>                     | A cable with one or more pairs of twisted insulated copper conductors bound in a single plastic sheath.                                                                                                                                                                        |
| <b>User Datagram Protocol (UDP)</b>                      | In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.                                                                                                                         |
| <b>user-based security model (USM)</b>                   | A security model that uses a defined set of user identities for authorized users on a particular Simple Network Management Protocol (SNMP) engine.                                                                                                                             |
| <b>view-based access control model (VACM)</b>            | Provides context, group access, and group security levels based on a predefined subset of management information base (MIB) objects.                                                                                                                                           |
| <b>Virtual Link Aggregation Control Protocol (VLACP)</b> | Virtual Link Aggregation Control Protocol (VLACP) is a Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.                                                                                                               |
| <b>Virtual Local Area Network (VLAN)</b>                 | A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.                                                                                    |
| <b>virtual router</b>                                    | An abstract object managed by the Virtual Router Redundancy Protocol (VRRP) that acts as a default router for hosts on a shared LAN.                                                                                                                                           |
| <b>virtual router forwarding (VRF)</b>                   | Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.                |
| <b>Virtual Router Redundancy Protocol (VRRP)</b>         | A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place. |